PARTITION-OF-UNITY ADAPTERS FOR EFFICIENT LOCAL REFINEMENT OF IMPLICIT PDE REPRESENTATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Pretrained implicit PDE surrogates and neural operators can produce imperfect predictions on new instances, while dense interior target solutions are unavailable for adaptation. Weight-space adapters reduce trainable parameters but do not explicitly localize corrections. We introduce *Partition-of-Unity Adapters*, a coordinate-space correction layer for frozen PDE representations. Smooth overlapping windows localize low-dimensional Chebyshev correction packets, tying each trainable coefficient to a known spatial support and basis mode. Residual or physical-change signals select packets; coordinated and dense variants address broader or coupled responses. The adapter is fit from physics constraints and optional sparse measurements, without dense interior solution labels. Across frozen learned solvers, singular-error repair, and material/interface changes, POU adapters often improve predictions over frozen models and matched global coordinate corrections. In held-out FNO tests with sparse measurements, 288 coefficients achieve useful accuracy–time tradeoffs against validation-selected full fine-tuning. Selected corrections preserve the frozen field outside their supports. The experiments show how support geometry, basis choice, and correction capacity affect performance across different error structures.

1 INTRODUCTION

Learned PDE representations—including physics-informed neural networks (Raissi et al., 2019), neural operators (Lu et al., 2021; Li et al., 2021), and coordinate-based neural solvers (Kairanda et al., 2026)—typically encode solution fields in shared network weights. An individual parameter has no prescribed spatial footprint, so an update need not be confined to a known region of the domain. When a pretrained field is accurate over most of the domain but underresolves a boundary layer, shock, or localized material response, fine-tuning these shared weights can degrade already accurate regions, a spatial form of catastrophic forgetting. This motivates an adaptation interface that assigns correction capacity to explicit physical supports and makes the extent of each update inspectable.

Existing approaches that introduce spatial structure into PDE learning address a different problem. Domain decomposition PINNs (Jagtap & Karniadakis, 2020; Moseley et al., 2023) build spatial structure into the solution representation during training rather than providing a correction layer for an existing model. Mixture-of-experts methods (Hu et al., 2023; Bonfanti et al., 2025) optimize routing and model weights jointly from scratch. Low-rank adapters (Hu et al., 2022a) and bottleneck insertion modules (Houlsby et al., 2019) enable parameter-efficient fine-tuning of frozen models, but their parameters are spatially opaque: a learned low-rank update does not identify which physical region was corrected.

We ask whether partition-of-unity correction packets can improve a pretrained PDE prediction using physics residuals, initial and boundary conditions, and optional sparse measurements, while making the correction spatially inspectable. We investigate this question across three settings: prediction errors in frozen learned solvers, localized errors near singularities, and material or interface changes that induce nonlocal solution responses.

A *Partition-of-Unity Adapter* comprises local *correction packets*: smooth POU windows multiply low-dimensional basis expansions, such as Chebyshev polynomials in normalized patch coordinates. The

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

windows localize the correction, tying each coefficient to a known physical support and basis mode. Deployment keeps the backbone g fixed and fits the selected packets' coefficients. Internal backbone weights are unnecessary, but physics-based fitting requires field evaluations and the derivatives or discrete operator used by the residual.

Spatial supports allow packets to be scored and activated by region. This gives three modes: selected (residual-scored packets), coordinated (selected packets with a limited coordination channel), and dense (all packets). Weight-space adapters do not directly provide this spatial control.

The main contributions are:

- **Partition-of-Unity Adapters.** A coordinate-space correction layer for frozen PDE representations, with compact packets attached to known physical regions and spatially inspectable coefficients.
- **Residual-guided adaptation of frozen models.** Fitting from physics constraints and optional sparse observations, with exact preservation outside selected supports. Matched held-out FNO tests quantify accuracy–time tradeoffs against validation-selected full fine-tuning.
- **Empirical evaluation across error structures.** Frozen-model adaptation, singular-error repair, and material/interface changes with nonlocal responses, compared with global controls and fine-tuning. Ablations identify when support geometry, fitting signals, problem-aware bases, coordination, or dense capacity improve correction quality.

2 Related work

Learned PDE representations and refinement. PINNs (Raissi et al., 2019) and neural operators such as DeepONet and FNO (Lu et al., 2021; Li et al., 2021) provide learned solution representations, including pretrained multi-task backbones such as Poseidon (Herde et al., 2024). Kernel models offer an alternative with spatially located basis functions. Refinement also improves learned representations: Lino & Thuerey (2026) generate fluid fields through coarse-to-fine sampling, while Xiong et al. (2026) refine implicit representations offline for fast queries. Our focus is deployment-time correction of an existing frozen PDE field using physics constraints and optional sparse measurements, with correction capacity assigned to explicit spatial supports.

Spatial structure in PDE learning. Domain-decomposition and mixture-of-experts PINNs (Jagtap & Karniadakis, 2020; Hu et al., 2023; Bonfanti et al., 2025) build spatial structure into the solution representation, training subdomain networks or learned routing together with the solution. XPINN analysis shows that decomposition helps only when the local problem is actually simpler, so that the reduction in subdomain complexity outweighs the loss of data per subdomain (Hu et al., 2022b). FBPINNs (Moseley et al., 2023) are the closest architectural relative: they sum subdomain networks confined by smooth overlapping windows, normalize inputs separately in each subdomain, and can train some subdomains while keeping others fixed. Their Schwarz interpretation adds additive and multiplicative subdomain updates and a coarse correction (Dolean et al., 2023). Interface-aware neural solvers similarly rebuild the representation around known material jumps (He et al., 2026). In all of these methods, the spatially structured components are the solution representation and are trained as part of it. POU adapters instead add a windowed correction, linear in its coefficients and optionally conditioned on the frozen field, to an existing prediction from an arbitrary backbone, and fit only these coefficients at deployment from physics constraints and optional sparse measurements.

Classical partition of unity and enrichment. Partition of unity, PUFEM, meshfree POU, and RBF-PU methods are established numerical-analysis tools for local enrichment of approximation spaces (Melenk & Babuška, 1996; Mir & Mirzaei, 2023). Recent variational enrichment work extends the same idea inside adaptive finite-element solvers (Dixon, 2026). This paper does not claim novelty for partition of unity itself. Classical POU enlarges a solver trial space; here the same local-window scaffold is attached as a compact coordinate-space adapter to a pretrained learned surrogate, and it is compared with ML adaptation baselines rather than as a replacement for adaptive FEM.

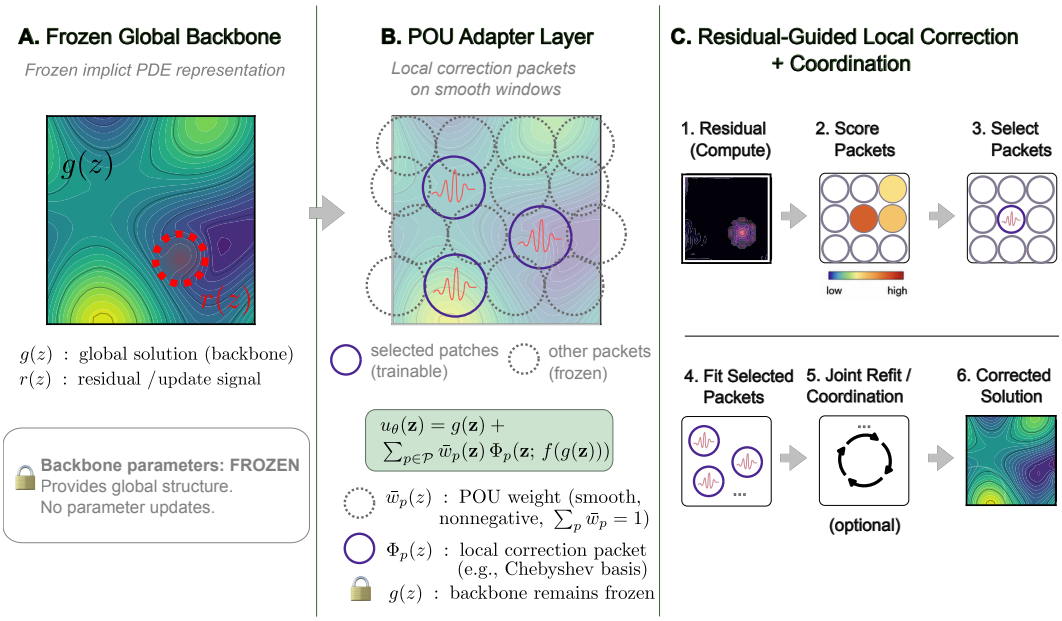


Figure 1: **Partition-of-unity adapter workflow.** Smooth overlapping POU windows attach local correction packets to a frozen backbone. Residual signals score candidate supports and select the correction packets to fit; an optional coordination step refines corrections globally without updating backbone parameters.

Parameter-efficient adaptation. LoRA (Hu et al., 2022a) and bottleneck adapters (Houlsby et al., 2019) reduce the number of trainable parameters when adapting pretrained models, but their updates are parameterized within the network. POU adapters instead fit an external coordinate-space correction, tying each coefficient to a prescribed physical support and local basis mode.

3 Method

The method is a coordinate-space adapter for an already trained PDE surrogate. Let g_{in} denote the fixed field available at adaptation time: it may be the prediction of a pretrained surrogate on a new instance, a corrupted numerical field, or a solution computed before a material change[1]. We write $\mathbf{z} = \mathbf{x}$ for steady problems and $\mathbf{z} = (\mathbf{x}, t)$ for time-dependent problems, where $\mathbf{x} \in \Omega$ is the spatial coordinate. The construction below uses spatial windows, with time included in the local packet basis. A POU adapter adds a trainable correction field A_θ outside the backbone:

$$u_\theta(\mathbf{z}) = g_{\mathrm{in}}(\mathbf{z}) + A_\theta(\mathbf{z}), \qquad A_\theta(\mathbf{z}) = \sum_{p \in \mathcal{P}} \bar{w}_p(\mathbf{x})\, \Phi_p(\mathbf{z};\, \eta(g_{\mathrm{in}}(\mathbf{z}))) \,. \tag{1}$$

Here $\mathcal{P}$ indexes a fixed family of overlapping spatial supports, $\bar{w}_p$ are normalized POU weights, and Φ_p is the local correction packet attached to support p. The trainable parameters are only the packet coefficients in θ; the backbone that produced g_{in} is frozen unless a comparison is explicitly labeled as a fine-tuning baseline. Figure 1 summarizes the workflow.

3.1 Adapter Parameterization

The adapter is defined by two ingredients: smooth POU windows that localize the correction, and a small local basis inside each window. Let $\mathcal{P}$ be a fixed family of overlapping supports covering the domain Ω. Each support $p \in \mathcal{P}$ has center c_p, radius r_p, and window $w_p(\mathbf{x}) = \varphi(\rho_p(\mathbf{x}))$ with $\rho_p(\mathbf{x}) = \|(\mathbf{x} - c_p)/r_p\|_2$, where $\varphi \geq 0$ is positive for $\rho < 1$, vanishes for $\rho \geq 1$, and yields windows

[1]The input field remains fixed during adaptation; updated material coefficients enter the PDE residual.

that are at least C^2. A standard choice is

$$w_p(\mathbf{x}) = \left[1 - \rho_p(\mathbf{x})^2\right]_+^3, \tag{2}$$

where $[\cdot]_+ = \max(\cdot, 0)$; it is C^2 because it and its first two derivatives vanish at $\rho_p = 1$. Anisotropic or transport-aligned supports replace ρ_p by an affine or characteristic local coordinate; in all cases the support geometry is fixed before fitting. Appendix A.3 lists the profile used for each benchmark. Under the coverage condition $\sum_{q\in\mathcal{P}} w_q(\mathbf{x}) > 0$, the normalized POU weights are

$$\bar{w}_p(\mathbf{x}) = \frac{w_p(\mathbf{x})}{\sum_{q\in\mathcal{P}} w_q(\mathbf{x})}. \tag{3}$$

so that $\sum_{p\in\mathcal{P}} \bar{w}_p(\mathbf{x}) = 1$ on covered points. Normalization makes the weights a partition of unity; exact preservation outside the active supports follows from compact support alone.

Each packet is evaluated in normalized local coordinates,

$$\boldsymbol{\xi}_p(\mathbf{x}) = \frac{\mathbf{x} - c_p}{r_p}. \tag{4}$$

For time-dependent problems on $[t_0, t_1]$ we optionally include the global normalized time coordinate

$$\tau(t) = 2\frac{t - t_0}{t_1 - t_0} - 1,$$

and omit this argument for steady problems. The default packet is a finite Chebyshev expansion modulated by fixed features of the frozen field,

$$\Phi_p(\mathbf{z}; \eta(g_{\text{in}}(\mathbf{z}))) = \sum_{m=1}^{M}\sum_{\ell=1}^{L} a_{p,m,\ell}\, \eta_\ell(g_{\text{in}}(\mathbf{z}))\, \psi_m(\boldsymbol{\xi}_p(\mathbf{x}), \tau(t)), \tag{5}$$

where $a_{p,m,\ell} \in \mathbb{R}$ are the trainable coefficients. Here, $\{\psi_m\}_{m=1}^M$ denotes a family of M tensor-product Chebyshev polynomials of the first kind up to a maximum degree d (so that each $\psi_m(\boldsymbol{\xi}) = \prod_{j=1}^D T_{i_j}(\xi_j)$ with local coordinates $\boldsymbol{\xi} \in [-1, 1]^D$; D counts the spatial coordinates and, when present, τ. Unless a total-degree bound is stated, $0 \le i_j \le d$ applies per coordinate, giving $M = (d+1)^D$), and $\eta(g_{\text{in}}(\mathbf{z})) \in \mathbb{R}^L$ is a fixed feature vector of length $L \in \{1, 2, 3\}$:

$$\eta(g_{\text{in}}(\mathbf{z})) \in \{[1],\ [1,\ g_{\text{in}}(\mathbf{z})],\ [1,\ g_{\text{in}}(\mathbf{z}),\ g_{\text{in}}(\mathbf{z})^2]\}. \tag{6}$$

For scalar fields, K active packets have $K \cdot M \cdot L$ coefficients, before any enrichment or coordination coefficients. Table 7, alongside the scaffold specification, lists the degree rule, features, and coefficient counts used in the experiments. For vector-valued fields, the polynomial features are applied componentwise unless stated otherwise. This factorization separates local coordinate dependence from the frozen solution branch: ψ_m determines the shape of the correction inside the packet, while $\eta(g_{\text{in}})$ lets the packet condition on what the backbone already predicts. Note that g_{in} therefore enters twice: as the fixed additive term in Eq. 1, which the adapter is never responsible for reproducing, and again here as an optional conditioning feature; setting $\eta = [1]$ removes the second role while leaving the additive residual structure intact.

Classical second-order residuals require the frozen field and the complete packet factors, including conditioning features, to have the required derivatives. Smooth windows alone do not regularize the frozen field or singular enrichments. Grid-valued fields use discrete residuals; singular-repair variants use the exclusions, energy formulations, or weak constraints specified in Appendix B.2.

For reference, the complete correction field is

$$A_{\theta_A}(\mathbf{z}) = \sum_{p\in\mathcal{P}} \bar{w}_p(\mathbf{x})\, \Phi_p(\mathbf{z}; \eta(g_{\text{in}}(\mathbf{z}))), \qquad \theta_A = \{a_{p,m,\ell}\}. \tag{7}$$

3.2 Fitting From Deployment Signals

The main deployment setting fits θ from information available at adaptation time. Dense interior target fields are excluded from the fitting objective and used only for evaluation. Residual evaluation

differentiates both the frozen field and the complete windowed, feature-conditioned correction, using analytic derivatives, automatic differentiation, or finite-difference stencils as specified per benchmark in Appendix A.6. We use a single adaptation objective family, of which any given deployment setting activates only the terms its available information supports,

$$\begin{aligned}\min_{\theta\in\Theta}\ \mathcal{L}_{\mathrm{adapt}}(\theta) = \lambda_{\mathrm{phys}}\,\mathcal{L}_{\mathrm{phys}}(u_\theta;\mathcal{P}_{\mathrm{new}}) + \lambda_{\mathrm{sens}}\,\|M_{\mathrm{sens}}u_\theta - y_{\mathrm{sens}}\|_2^2 \\ + \lambda_{\mathrm{anchor}}\,\|M_{\mathrm{anchor}}(u_\theta - g_{\mathrm{in}})\|_2^2 + \lambda_{\mathrm{ridge}}\|\theta\|_2^2.\end{aligned} \tag{8}$$

Here $u_\theta = g_{\mathrm{in}} + A_\theta$. The core term $\mathcal{L}_{\mathrm{phys}}$ is the physics-informed defect of the new deployment problem $\mathcal{P}_{\mathrm{new}}$: it combines the PDE residual with the available initial, boundary, material, or interface constraints. When prescribed boundary values change, a fixed lifting function first accounts for the difference between the previous and updated boundary data; the adapter then corrects the lifted field. Sparse measurements enter only through M_{sens} when sensors are available. The preservation term encodes a prior that the frozen field should change little on M_{anchor}, chosen from deployment information. Without such a prior, $\lambda_{\mathrm{anchor}} = 0$. An unchanged material coefficient alone does not establish that the solution should remain unchanged. The ridge term penalizes large packet coefficients, regularizing the correction toward zero. Exact preservation outside active supports follows from the compact windows when no global coordination term is present. Coordinated relaxation variants may additionally penalize the correction norm; Appendix A.8 gives the corresponding least-squares instantiation. Stagewise line searches and proximal refits can add implementation-specific acceptance or preservation penalties. Appendix A.9 specifies residual-block normalization, regularization, and solver settings.

Because the correction is linear in its coefficients, linear PDE and constraint operators yield a regularized least-squares problem. For nonlinear PDEs, coefficients can be fitted by linearizing the residual around the current field, optionally repeating the update with safeguarded Gauss–Newton steps, or by directly optimizing the nonlinear objective. Appendix A.7 gives the linearized formulation; Appendices A.6 and A.9 specify the benchmark solvers and budgets. Supervised residual fitting against a dense reference field is not part of the deployment method; it appears only as a diagnostic capacity test in the appendix.

3.3 Operating Modes

Equation 1 defines a single POU adapter family. At deployment, we instantiate this family at three levels of spatial capacity, summarized in Table 1. The modes differ only in which adapter coefficients are active and whether a limited coordination channel is included; the pretrained backbone remains fixed.

Table 1: **POU adapter operating modes.** POU-S is the default compact adapter; POU-C keeps the selected packets but adds limited coordination for coupled responses; POU-D activates the full scaffold for dense adaptation when errors are not localized.

Mode	Active correction space	Deployment regime
POU-S	top-K packets by residual score	localized residual, trace, or interface defects
POU-C	selected packets plus a low-order relaxation, coarse-to-fine, or subdomain-coupled channel	broad or coupled responses needing packet coordination
POU-D	all packets in the fixed scaffold	distributed defects; capacity reference for compact modes

Formally, let $S_K \subseteq \mathcal{P}$ be the selected support set from Table 1: POU-S sets coefficients outside S_K to zero, POU-C keeps S_K but adds its coordination channel, and POU-D sets $S_K = \mathcal{P}$. The residual scoring in Figure 1 ranks packets by their windowed residual energy $s_p = \sum_{\mathbf{z}\in Q}\left(\bar{w}_p(\mathbf{x})\, r[g_{\mathrm{in}}](\mathbf{z})\right)^2$ over the interior collocation set Q, where $r[g_{\mathrm{in}}]$ is the residual of the new problem, and activates the K largest (Appendix A.5). These three modes are the method variants used in the main text; support geometry, flux penalties, sensor-aware reselection, mean-field bases, hierarchy, LoRA, matched global corrections, and full fine-tuning are treated as ablations, objective choices, or baselines.

4 EXPERIMENTS

We evaluate POU adapters for refining frozen PDE representations. The input field g_{in} comes from a pretrained coordinate model, a neural operator, or a previously computed solution. Each POU variant keeps this field fixed and fits only adapter coefficients. Deployment fits use PDE residuals, initial or boundary conditions, interface conditions, preservation priors constructed from deployment information, and optional sparse measurements. Dense test fields are reserved for evaluation in these experiments. In the matched comparison, separate dense validation fields select optimization settings and budgets. Supervised diagnostics that fit adapters using dense reference fields are identified explicitly.

Normalization and fitting procedures vary across benchmarks. Only the matched FNO comparison uses an identical normalized objective for coordinate corrections and full fine-tuning. Depending on the setting, we report solution and constraint errors, outside-region preservation, trainable parameter counts, and adaptation time. Table 2 uses an AMD EPYC 7742 CPU and NVIDIA RTX 4090 GPU (24GB); earlier workstation measurements are reported separately. Appendix A.9 specifies the fitting and timing protocols.

The section is organized by correction structure. We first test the main deployment setting: adapting frozen learned solvers on held-out PDE instances. We then analyze localized and singular errors, followed by structured physical changes whose solution response is nonlocal. Finally, we report mechanism and capacity diagnostics, including multiscale tests, to isolate when locality, coordination, basis choice, and dense capacity matter.

4.1 ADAPTING FROZEN LEARNED SOLVERS

We evaluate frozen FNOs after 100 pretraining steps on Allen–Cahn, Burgers, and convection–diffusion–reaction (CDR), with 32 training, eight validation, and 12 test instances per PDE and three backbone seeds. Validation selects adaptation budgets before testing. Table 2 compares POU-S, Global Cheb., and full fine-tuning (FT) under identical objectives and observations. Global Cheb. uses an unwindowed, domain-wide correction with POU-S's features $[1, g, g^2]$, nominal 288-coefficient budget, and safeguarded Gauss–Newton fitting. The paired 350-step comparison appears in Appendix C.1.

Table 2: **Held-out FNO adaptation after 100 pretraining steps.** Relative L_2 mean $\pm$ sample SD over 12 instances $\times$ three backbone seeds (crossed observations). Coordinate counts are nominal. $m = 0$ uses PDE/IC/BC; $m = 32$ adds sensors. Times are mean per-case median warm seconds, including setup. Parentheses show validation-selected GN/Adam iteration budgets; FT validation searches extend to 5600 steps. Protocol and supporting results: Appendices A.9 and C.1.

PDE	m	Frozen	Global Cheb.	POU-S (GN)	FT
Trainable parameters		0	288	288	241K
Allen–Cahn	0	0.0901 ± 0.0266	0.0222 ± 0.0207	0.0525 ± 0.0458	0.1148 ± 0.0254
Time (s)		0.00109	0.123 (2)	0.122 (2)	0.0581 (50)†
Allen–Cahn	32	0.0901 ± 0.0266	0.0117 ± 0.0073	0.0178 ± 0.0161	0.0490 ± 0.0518
Time (s)		0.00108	0.205 (4)	0.131 (2)	6.08 (5600)
Burgers	0	0.2594 ± 0.1276	0.0271 ± 0.0140	0.0313 ± 0.0161	0.0204 ± 0.0132
Time (s)		0.00109	0.200 (8)	0.204 (8)	3.06 (2800)
Burgers	32	0.2594 ± 0.1276	0.0237 ± 0.0124	0.0272 ± 0.0144	0.0177 ± 0.0104
Time (s)		0.00109	0.219 (8)	0.222 (8)	3.06 (2800)
CDR	0	0.1962 ± 0.0393	0.1020 ± 0.0659	0.1098 ± 0.0649	0.0197 ± 0.0154
Time (s)		0.00113	0.0883 (1)	0.0872 (1)	6.18 (5600)
CDR	32	0.1962 ± 0.0393	0.0896 ± 0.0677	0.0964 ± 0.0656	0.0113 ± 0.0081
Time (s)		0.00112	0.0932 (1)	0.0924 (1)	3.10 (2800)

† For Allen–Cahn with $m = 0$, 50 steps minimize validation error among positive FT budgets; allowing zero steps would select the frozen model.

POU reduces mean frozen-field error by 42–90% across these conditions in 0.087–0.222 seconds. With 32 sensors, Allen–Cahn error is 0.0178 versus FT's 0.0490, at 0.131 versus 6.08 seconds.

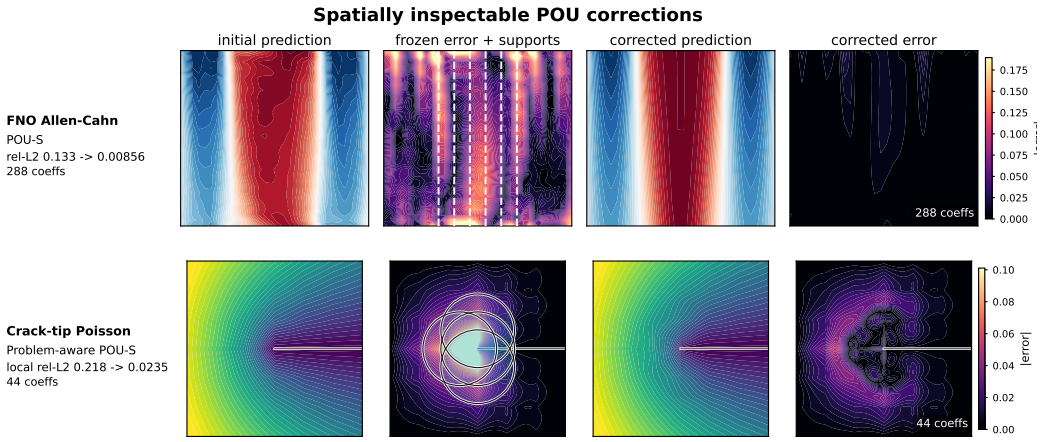


Figure 2: **Spatially inspectable POU corrections.** Top: selected POU-S packets adapt a frozen FNO Allen–Cahn prediction on one instance (horizontal axis x, vertical axis t), using a separate backbone configuration (80 pretraining steps, width 16); dashed lines mark the selected packet centers, whose supports (radius 0.72) jointly cover the domain. Bottom: geometry-aware POU-S repairs a crack-tip Poisson solution near the singular region (panels in the (x, y) plane); circles mark the selected compact supports, outside which the frozen field is unchanged. Dense reference fields are used only to compute the displayed errors and evaluation metrics, not to fit the adapters.

FT is more accurate on Burgers and CDR; Global Cheb. is more accurate than POU throughout. Sensor-free FT does not improve the Allen–Cahn frozen field. These results show problem-dependent accuracy–cost tradeoffs; paired outcomes and constraints appear in Appendices C.1 and C.3.

Selected supports cover 81.8–100% of the domain here; the following sections test preservation around localized defects. Appendix C.4 reports fixed-budget comparisons with one-shot FNO fitting, characteristic-aligned CDR supports, and same-instance FF-MLP refinement. The FF-MLP fit activates all packets and is labelled selection-initialized POU-D; its validation-selected FT comparison appears in Appendix C.2.

4.2 Error Geometry: Localized and Singular Errors

We next test whether local corrections can repair underresolved singular features while preserving accuracy elsewhere. We consider three Poisson benchmarks: an L-shaped domain with a re-entrant corner, a slit domain with a crack-tip singularity, and a domain containing a point source. In each case, the input is a coarse numerical solution interpolated onto a finer grid, and adaptation uses physics constraints alongside 32 sparse measurements (Appendix B.2). Table 3 compares correction quality within the defect against outside-region preservation.

For the corner and crack-tip benchmarks, global Chebyshev corrections reduce local error but increase outside-region error. Selected POU-S packets improve local accuracy while maintaining or reducing outside-region error. The selected supports leave 62–75% of each domain unchanged; most of each outside region (84–96%) lies beyond them, and the remainder lies in the support fringe, where the correction can still change the error (Appendix A.3). On the crack-tip problem, for example, the global correction increases outside error from 0.0135 to 0.0335, whereas POU-S retains 0.0135 at the reported precision. Enriching the local basis with the known crack-tip singular function further reduces local error from 0.0278 to 0.0167. Figure 2 (bottom) illustrates the selected supports and resulting correction. These results support localized correction for geometric singularities, with additional gains from a basis that incorporates the known singular structure.

The point-source case exposes a limitation of smooth local packets. Standard POU-S reduces local error but increases outside-region error from 0.0202 to 0.0406. About 15% of this outside evaluation region lies within the selected packet supports, where the correction can change the error; the strict

Table 3: **Local-error structure.** Local repair is successful only if it reduces defect-region error without increasing outside-region error. Entries report local rel-L2, with outside-region rel-L2 in parentheses, for the $m = 32$ sparse-sensor setting, using the separate evaluation masks in Appendix B.2. Geometry/source-aware POU uses the same selected-packet family with singular or weak-physics enrichment. The Global column is a matched-budget nonlocal Chebyshev correction, using the same parameter count as POU-S. "Params/time" gives the range across Global, POU-S, and Geometry/source-aware POU; times are archived fit-stage measurements excluding common field construction.

Error structure	Frozen	Global	POU-S	Geometry/source-aware POU	Params/time
L-shaped corner	0.1256 (0.0188)	0.0627 (0.0559)	0.0653 (0.0168)	**0.0601** (0.0163)	40–44/0.18–0.27s
Crack tip	0.1591 (0.0135)	0.0782 (0.0335)	0.0278 (0.0135)	**0.0167** (0.0135)	40–44/0.06–0.07s
Point source	0.0823 (0.0202)	0.0811 (0.0723)	0.0564 (0.0406)	**0.0298** (0.0246)	40–44/0.06–0.07s

Table 4: **Adaptation with nonlocal responses.** Entries report relative L_2 error against the manufactured target. The κ rows fit bulk PDE and outer-boundary constraints only; the circle/sunflower rows also fit interface conditions. Constraint errors appear in Table 21. "Subdomain POU-C" uses separate packet sets coupled through interface constraints. Dashes indicate variants not reported in this table. "Cost" reports trainable parameter counts and archived fit-stage times, excluding common field construction and diagnostics.

Problem setting	Frozen	Global	POU-S	POU-C	POU-D	Subdomain POU-C	Cost
κ, $k_R = 2$	0.1065	0.1015	0.0900	**0.0565**	0.0624	–	144–576/0.01–0.10s
κ, $k_R = 5$	0.1184	0.1152	0.1128	**0.1015**	0.1059	–	144–576/0.01–0.10s
κ, $k_R = 10$	0.1284	0.1239	0.1192	0.1264	**0.1184**	–	144–576/0.01–0.10s
Circle interface	0.587	0.560	–	**0.433**	–	0.479	166–182/0.05–0.24s
Sunflower interface	0.878	1.683	–	3.582	–	**0.694**	166–182/0.05–0.24s

support exterior remains unchanged (Appendix A.3). The source-aware variant adds Green-function enrichment, evaluates residuals in an annulus around the source, and imposes an integrated flux constraint on a surrounding circle to encode the source strength (Appendix B.2). Together, these additions reduce local error to 0.0298 and outside-region error to 0.0246. The latter remains above the frozen-field error, indicating a remaining preservation tradeoff. These results motivate matching the local basis and fitting constraints to the singular structure of the problem.

4.3 Error Geometry: Structured Changes With Nonlocal Response

Changes in material properties can require corrections beyond the changed region. We examine two steady-diffusion settings with spatially varying conductivity $\kappa(\mathbf{x})$ (Appendix B.3). The variable-conductivity diagnostic starts from a prescribed smooth field and fits bulk PDE residuals and outer-boundary conditions without interface constraints. Its manufactured operator and forcing vary with the conductivity contrast $k_R = \kappa_{\text{right}}/\kappa_{\text{left}}$. The circle and sunflower benchmarks instead adapt a previous material-instance solution, lifted for changed boundary data where necessary, using interface value and flux constraints. Table 4 reports relative solution errors.

We compare selected packets (POU-S), selected packets plus a low-order domain-wide correction (POU-C), and the full packet scaffold (POU-D). In the variable-conductivity diagnostic, POU-S leaves 34–57% of the domain unchanged; POU-C and POU-D act throughout.

For $k_R = 2$ and $k_R = 5$, POU-C achieves the lowest reported errors, 0.0565 and 0.1015, respectively, improving on POU-S and the global coordinate correction. At $k_R = 10$, however, POU-C produces higher error than POU-S (0.1264 versus 0.1192), while POU-D achieves 0.1184. Coordination therefore improves solution accuracy for some contrasts, but not uniformly. Interface flux continuity is not enforced in this diagnostic and cannot be inferred from these errors. Fit-stage times are 0.01–0.10 seconds (medians of three repeats), excluding common field construction and diagnostics.

For the circular interface benchmark, with a continuous solution, standard POU-C achieves the lowest reported error (0.433). For the sunflower benchmark, with a prescribed solution jump, standard POU-C increases error from 0.878 to 3.582. Subdomain POU-C couples separate inside and outside packet sets through interface constraints and reduces the sunflower error to 0.694, improving on the frozen field while leaving substantial solution error. Boundary and interface errors are reported separately in Table 21.

4.4 Mechanisms and Capacity

The preceding sections report deployment performance; here we use appendix diagnostics to isolate the main design factors. Capacity is controlled by the active-packet budget K and the local basis degree. A damaged-repair budget sweep shows that the local-error conclusion is not tied to the 40-coefficient setting: crack-tip POU-S improves over matched global Chebyshev at 20, 40, and 80 coefficients, while source-aware enrichment yields the lowest point-source local error at all three budgets (Table 26).

The multiscale diagnostic tests whether local packets can recover fine structure smoothed by the frozen coarse field. On a 1D elliptic benchmark with a rapidly oscillating coefficient and a supervised fit to the dense target, POU packets reduce error by up to 39% relative to the frozen coarse field (Appendix D.4). This supervised comparison is a capacity diagnostic. Singular-repair errors are not monotone in the coefficient budget (Table 26). More capacity alone does not ensure an accurate correction; fitting signals and problem-aware correction spaces also matter.

5 Discussion, Limitations, and Future Work

The experiments highlight two practical considerations for adapting frozen PDE representations: the spatial structure of the required correction and the information available to constrain it. Selected packets (POU-S) are effective when their supports and bases match the error structure. Broader responses can benefit from coordination (POU-C) or additional capacity, though these extensions do not consistently outperform selected packets. Physics residuals alone may also provide insufficient guidance for accurate reconstruction under the evaluated fitting procedures. Sparse measurements, preservation priors, and problem-aware bases or constraints can improve correction quality, but their benefits depend on the problem.

Scope and limitations. Support geometries, packet budgets, and basis degrees are prescribed; automatic selection remains open. Broader validation across PDE families and held-out instances is needed: the primary FNO comparison covers three PDEs and three backbone seeds. Validation-selected, budget-limited FT searches do not establish convergence, and timing comparisons do not establish superiority over fully converged fine-tuning or adaptive numerical solvers. Residual rows are assembled densely despite the small coefficient count. Compact-support assembly could reduce cost while retaining terms required for nonlocal coupling.

Future work. Frozen-field structure, residual indicators, and error estimators could guide support construction, basis enrichment, and coordination. Physics-guided inpainting could propose corrections in uncertain regions, constrained by physics and trusted surrounding values. Distillation could transfer localized POU corrections into an operator backbone for subsequent deployment.

6 Conclusion

Partition-of-unity adapters refine frozen PDE representations from physics constraints and optional sparse measurements, tying each coefficient to a spatial support and basis mode. Selected corrections preserve the field outside their supports; coordinated and dense variants extend the correction space. The experiments support targeted refinement, with effectiveness governed by response geometry, correction capacity, and fitting information. Compact adapters may also support low-latency correction of surrogates deployed on resource-constrained devices.

Reproducibility Statement

The appendix specifies the benchmark equations, adaptation signals, baselines, trainable-parameter budgets, and diagnostic studies used to support the main tables. Dense test fields are used only for evaluation unless a row is explicitly identified as an oracle or supervised diagnostic; separate dense validation fields select the reported optimization recipes.

We will release code, including the optimizer implementations, training scripts and the configuration files of all experiments, together with the final version of the paper.

Ethics Statement

This work studies numerical methods for adapting PDE surrogates and does not involve human subjects or personal data. As with any surrogate modeling method, deployment in safety-critical simulation pipelines would require independent validation against trusted solvers and domain-specific acceptance criteria.

AI Use Statement

We used LLMs to help implement the method and experiment code (including the generators of the manufactured datasets) and the plotting scripts, to launch and monitor runs, to aggregate results, to check the implementation against the manuscript, to provide feedback on the manuscript, to search the literature, and to improve the wording of the text. LLMs were not used to develop theoretical models, conceptual frameworks, mathematical claims, or hypotheses, to write proofs, or to translate. The authors made all decisions on methodology, experimental design, and interpretation, verified every issue raised by LLM-assisted checks against the code and results, and checked every cited work. All datasets were produced by numerical solvers or manufactured formulas, without LLM-assisted cleaning or reformatting. The authors reviewed all LLM-generated code and text and take full responsibility for the content of this work.

References

Andrea Bonfanti, Ismael Medina, Roman List, Björn Staeves, Roberto Santana, and Marco Ellero. PINN balls: Scaling second-order methods for PINNs with domain decomposition and adaptive sampling. In *Advances in Neural Information Processing Systems*, 2025.

Matthew Francis Dixon. Variational enrichment of adaptive finite element spaces beyond polynomials. *arXiv preprint arXiv:2608.11764*, 2026. doi: 10.48550/arXiv.2608.11764.

Victorita Dolean, Alexander Heinlein, Siddhartha Mishra, and Ben Moseley. Finite basis physics-informed neural networks as a Schwarz domain decomposition method, 2023. URL `https://arxiv.org/abs/2211.05560`.

Tingni He, Kenan Chen, Beibei Wang, and Xinxiang Li. Deep learning-based hard constraint method for solving elliptic interface problems. *Engineering Analysis with Boundary Elements*, 191:106920, 2026. doi: 10.1016/j.enganabound.2026.106920.

Maximilian Herde, Bogdan Raonic, Tobias Rohner, Roger Kaeppeli, Roberto Molinaro, Emmanuel de Bezenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for PDEs. *arXiv preprint arXiv:2405.19101*, 2024.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning*, pp. 2790–2799, 2019.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022a.

Zheyuan Hu, Ameya D. Jagtap, George Em Karniadakis, and Kenji Kawaguchi. When do extended physics-informed neural networks (XPINNs) improve generalization? *SIAM Journal on Scientific Computing*, 44(5):A3158–A3182, 2022b.

Zheyuan Hu, Ameya D. Jagtap, George Em Karniadakis, and Kenji Kawaguchi. Augmented physics-informed neural networks (APINNs): A gating network-based soft domain decomposition methodology. *Engineering Applications of Artificial Intelligence*, 126:107183, 2023.

Ameya D. Jagtap and George Em Karniadakis. Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition framework for nonlinear partial differential equations. *Communications in Computational Physics*, 28(5):2002–2041, 2020.

Navami Kairanda, Shanthika Naik, Marc Habermann, Avinash Sharma, Christian Theobalt, and Vladislav Golyanik. ∂-grid: A neural differential equation solver with differentiable feature grids. *arXiv preprint arXiv:2601.10715*, 2026.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

Mario Lino and Nils Thuerey. One scale at a time: Scale-autoregressive modeling for fluid flow distributions. *arXiv preprint arXiv:2604.11403*, 2026.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqing Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.

J. M. Melenk and I. Babuška. The partition of unity finite element method: Basic theory and applications. *Computer Methods in Applied Mechanics and Engineering*, 139(1–4):289–314, 1996. doi: 10.1016/S0045-7825(96)01087-0. URL https://www.sam.math.ethz.ch/sam_reports/reports_final/reports1996/1996-01.pdf. Also available as SAM Research Report 96-01, ETH Zürich.

Reyhaneh Mir and Davoud Mirzaei. The D-RBF-PU method for solving surface PDEs. *Journal of Computational Physics*, 479:112001, 2023.

Ben Moseley, Andrew Markham, and Tarje Nissen-Meyer. Finite basis physics-informed neural networks (FBPINNs): a scalable domain decomposition approach for solving differential equations. *Advances in Computational Mathematics*, 49(4), July 2023. ISSN 1572-9044. doi: 10.1007/s10444-023-10065-9. URL http://dx.doi.org/10.1007/s10444-023-10065-9.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

Tianyu Xiong, Skylar W. Wurster, and Han-Wei Shen. Refine now, query fast: A decoupled representation refinement paradigm for implicit neural fields. In *International Conference on Learning Representations*, 2026.

A ADAPTER CONSTRUCTION, FITTING, AND EVALUATION PROTOCOLS

Appendix A collects adapter construction, coefficient fitting, and evaluation protocols; Appendix B gives benchmark equations and settings in the order of the main experiments. Appendix C presents supporting deployment results, historical comparisons, and ablations; Appendix D groups controlled locality and capacity diagnostics. Appendix E contains the supplementary RBF construction and solve scaling diagnostic. Deployment results use adaptation-time signals; supervised and oracle diagnostics are identified separately.

A.1 EVALUATION METRICS AND OBSERVATION SAMPLING

Deployment fits use only adaptation-time signals, with dense target fields withheld except for supplied sparse measurements. Supervised residual fits are identified as locality or capacity diagnostics. The following conventions define the reported errors and observation sets.

Relative solution error is the primary comparison metric; Appendix C.3 reports companion constraint diagnostics for the main rows. Relative solution error and PDE residual RMSE are defined as

$$E_{\mathrm{rel}}(u_{\theta_A}) = \frac{\|u_{\theta_A} - u^\star\|_{2,\Omega_{\mathrm{eval}}}}{\|u^\star\|_{2,\Omega_{\mathrm{eval}}}}, \qquad R_{\mathrm{pde}}(u_{\theta_A}) = \left(|\Omega_{\mathrm{int}}|^{-1}\|\mathcal{R}[u_{\theta_A}]\|^2_{2,\Omega_{\mathrm{int}}}\right)^{1/2}. \tag{9}$$

The norms are unweighted discrete norms over the stated evaluation sets,

$$\|v\|^2_{2,Q} = \sum_{\mathbf{z}\in Q} |v(\mathbf{z})|^2, \qquad |Q| = \text{number of evaluation points in } Q. \tag{10}$$

For sparse FNO results, Ω_{eval} is the full space–time grid, including sensor locations and boundary/initial points. Table 27 instead uses the outer-body mask (including internal void locations) and the valid solid-domain mask of configuration B; both retain sensor locations. The reference is zero in void cells, while predictions are not masked there, so the outer-body score also penalizes nonzero values inside voids; the frozen field keeps the configuration-A temperatures in cells that become void. The solid-domain error of configuration B is the physically meaningful metric. In the FNO rows, sensor locations are drawn once per run, uniformly without replacement from the interior grid points (excluding the initial time and the spatial boundary); they are nested across sensor counts, shared by all methods, and noise-free. Unobserved errors for the matched multi-instance comparison appear in Table 18. In the singular benchmarks, sensors are drawn uniformly from grid points inside the selected supports (for the point source, from the $8m$ points nearest the source), that is, preferentially near the defect. Here, held-out refers to learned-solver test instances, not to a sensor-excluded evaluation mask.

Residual RMSE is an equation-satisfaction diagnostic on Ω_{int}, not a disjoint-point validation score. Residual evaluation points can also enter coefficient fitting; the operators and grids are specified in Appendix A.6.

Localized/singular-error results in Table 3 use separate local and outside evaluation masks, D_{loc} and D_{out}, specified in Appendix B.2. Table 26 instead pairs local with full-domain error. The regional errors are

$$E_{\mathrm{loc}} = \frac{\|u_{\theta_A} - u_{\mathrm{ref}}\|_{2,D_{\mathrm{loc}}}}{\|u_{\mathrm{ref}}\|_{2,D_{\mathrm{loc}}}}, \qquad E_{\mathrm{out}} = \frac{\|u_{\theta_A} - u_{\mathrm{ref}}\|_{2,D_{\mathrm{out}}}}{\|u_{\mathrm{ref}}\|_{2,D_{\mathrm{out}}}}. \tag{11}$$

These masks need not partition the evaluation grid. In particular, D_{out} can include points inside active window supports; exact preservation of the frozen field applies only where all active windows vanish. The saved interface diagnostics measure mismatch to the prescribed jumps in Appendix B.3; Table 4 reports solution error, not these diagnostics:

$$J_u = \|[\![u_{\theta_A}]\!] - q_u\|_{\Gamma,\mathrm{rms}}, \qquad J_q = \|[\![\alpha\nabla u_{\theta_A}\cdot n]\!] - q_\alpha\|_{\Gamma,\mathrm{rms}}, \tag{12}$$

where α is the updated material coefficient, $[\![v]\!] = v_{\mathrm{out}} - v_{\mathrm{in}}$, and n points from the inner to the outer subdomain. We use $\|r\|_{\Gamma,\mathrm{rms}} = (N_\Gamma^{-1}\sum_{j=1}^{N_\Gamma}|r(\mathbf{x}_j)|^2)^{1/2}$ over N_Γ interface samples; homogeneous transmission sets $q_u = q_\alpha = 0$. The circle and sunflower benchmarks use 128 samples uniformly spaced in polar angle, with traces approximated at offset points on either side.

A.2 COORDINATE-SPACE CONTROLS AND COMPARISON SCOPE

Each benchmark compares the applicable frozen, coordinate-space, weight-space, and POU variants in Section 4. Full fine-tuning updates the backbone from the same deployment signals; LoRA adds weight-space updates. POU-S, POU-C, and POU-D use the correction spaces in Table 1.

Global coordinate-space controls compare POU packets with nonlocal correction spaces under the matching conditions stated for each experiment. Let g be the frozen backbone and let $\{h_\ell(\mathbf{z})\}_{\ell=1}^{L}$ be a nonlocal coordinate basis, such as Fourier features, Chebyshev features, or frozen backbone feature columns. The global coordinate correction baseline uses

$$u_{\mathrm{global}}(\mathbf{z}) = g(\mathbf{z}) + \Delta_\beta(\mathbf{z}), \qquad \Delta_\beta(\mathbf{z}) = \sum_{\ell=1}^{L} \beta_\ell h_\ell(\mathbf{z}), \tag{13}$$

and fits only the coefficients β from the available supervised or physics-informed signal. The primary FNO comparison matches features, coefficient counts and the normalized objective; historical controls match only the stated budgets/signals and may differ in conditioning, regularization or solver. This keeps the backbone fixed, but the correction is nonlocal: the coefficients are not tied to individual physical patches.

The masked or topology-window global control gives this nonlocal correction the same spatial support information used by POU selection. Given a fixed mask $m(\mathbf{z}) \in [0, 1]$ from the selected update region or known topology-change region, it uses

$$u_{\mathrm{mask}}(\mathbf{z}) = g(\mathbf{z}) + m(\mathbf{z})\, \Delta_\beta(\mathbf{z}). \tag{14}$$

This is a support-informed control: it fits one global coordinate correction inside the supplied window rather than separate normalized POU packets in local coordinates. Access to a support mask is distinct from access to dense target errors for oracle packet selection. Table 29 reports a supervised diagnostic in which both representations receive the same selected-window mask, features and coefficient budget.

For the kernel-backbone local-damage experiment, the global RBF-feature delta uses the frozen RBF activations $\phi_i(\mathbf{z})$ as the nonlocal basis,

$$u_{\mathrm{rbf}\Delta}(\mathbf{z}) = g(\mathbf{z}) + \sum_{i=1}^{M} \beta_i \phi_i(\mathbf{z}). \tag{15}$$

Only the added output coefficients β_i are fit; the RBF centers, widths, rotations, and original weights remain frozen. Full fine-tuning updates the original backbone parameters directly. In Table 2, Global Cheb. denotes the conditioned nonlocal correction and FT denotes full fine-tuning. In the historical Table 23, Control denotes a coefficient-budget control for FF-MLP/FNO or standard spatial POU for the CDR geometry comparison. Tables 3 and 4 label their global coordinate controls "Global".

A.3 WINDOW PROFILES AND CANDIDATE SCAFFOLDS

Window profiles and candidate support geometries are prescribed for each benchmark using known domain or PDE structure where appropriate. Active packets are selected from adaptation signals, and geometry remains fixed during coefficient fitting. All windows use compactly supported radial profiles $\varphi(\rho_p)$ of the scaled distance in Section 3.1, fixed before fitting. Table 5 lists the profile, normalization, and local coordinate for each benchmark. Table 6 specifies the candidate supports for the main experiments and porous deployment; numerical normalization and coverage are described below.

Characteristic-aligned CDR packets. In the historical CDR geometry comparison, both the windows and the spatial Chebyshev factors use periodic distance on $[-1, 1]$ in the characteristic coordinate $\zeta = x - vt$, where v is the known PDE velocity; the temporal factors retain $\tau(t)$. Packet centers and radii are fixed in characteristic coordinates, so their spatial supports move with time. These runs evaluate the complete conditioned basis on the (x, t) grid and apply the same finite-difference residual operator used for the frozen field (Appendix A.6). The matched CDR comparison in Table 2 instead uses the default spatial windows and ordinary distance in x.

Table 5: **Window profiles by benchmark.** All radial profiles are compactly supported and at least C^2; normalized variants use the additive stabilization described below.

Benchmarks	Profile $\varphi(\rho)$	Norm.	Local coordinate
Frozen learned solvers (Tables 2, 23, 24; Fig. 2, top)	Eq. 2	yes	$(x - c_p)/r_p$; 12 centers on $[-1, 1]$, $r_p = 0.72$; periodic distance for the Burgers FF-MLP backbone
Historical CDR, characteristic-aligned packets	Eq. 2 in $\zeta = x - vt$, known v	yes	$(\zeta - c_p)/r_p$, periodic in ζ
Singular repair (Tables 3, 26; Fig. 2, bottom)	Wendland $(1-\rho)_+^4(4\rho+1)$	no	$(\mathbf{x} - c_p)/r_p$
Material, interface, porous, and heat updates; supporting diagnostics (Tables 4, 27, 30, 31, 32; Appendices D.1–D.2, D.4)	Eq. 2	yes	$(\mathbf{x} - c_p)/r_p$

Candidate scaffolds. Write $\mathcal{M}_n(a, b) = \{a + (i + \frac{1}{2})(b - a)/n\}_{i=0}^{n-1}$ for midpoint centers and $\mathcal{U}_n(a, b) = \{a + i(b - a)/(n - 1)\}_{i=0}^{n-1}$ for an endpoint-inclusive grid; $X^2 = X \times X$. Centers and radii below are in physical coordinates, with local coordinates $(\mathbf{x} - c_p)/r_p$ (or periodic displacement divided by r_p). Candidate sets are fixed before packet selection, and unions remove duplicate centers.

Table 6: **Candidate support scaffolds.** N counts candidate centers before selection, not active packets. Radii determine the overlap. The frozen-solver radius ablation additionally uses $r_p = 1/6$ and $1/3$ (Appendix C.6).

Benchmark	Candidate centers	N	r_p
Frozen learned solvers	$\mathcal{M}_{12}(-1, 1)$ in x	12	0.72
Characteristic CDR	$\mathcal{M}_{12}(-1, 1)$ in $\zeta = x - vt$	12	0.72
L-shaped corner	$\mathcal{U}_7(-0.82, 0.82)^2$, removing $x > 0, y < 0$ and $(0, 0)$; add E_L	45	0.42
Crack tip	$\mathcal{U}_9(-0.82, 0.82)^2 \cup E_C$	87	0.42
Point source	$\mathcal{U}_9(0.12, 0.88)^2 \cup \{0.38, 0.5, 0.62\}^2$	89	0.26
Variable conductivity	$\mathcal{M}_4(0, 1)^2$	16	0.31
Circle/sunflower interfaces	$\mathcal{M}_6(-1, 1)^2$; same centers for each subdomain copy	36	0.36
Porous topology update	$\mathcal{M}_6(0, 1)^2$ on the bounding square	36	0.17

The extra singular-region centers are

$$E_L = \{(-0.18, 0.18), (-0.18, -0.18), (0.18, 0.18), (-0.34, 0), (0, 0.34), (0, -0.34)\},$$
$$E_C = \{(-0.18, 0), (0, 0.18), (0, -0.18), (0.18, 0.18), (0.18, -0.18), (0.34, 0)\}.$$

The Burgers FF-MLP windows and characteristic CDR windows use period-two distance; the other frozen-solver windows, including the standard CDR control, use ordinary distance. Two-dimensional windows use Euclidean distance, with no periodic images or added boundary centers. Singular-repair packet columns are zeroed outside the L-shaped domain, on the crack, or on the point-source square boundary, respectively. Subdomain interface variants use independent coefficient sets restricted to the inside/outside masks, with normalization over the same full 36-center scaffold. Porous residuals are masked to valid solid cells while retaining the full bounding-square scaffold.

Numerical normalization and coverage. Normalized implementations use an additive denominator stabilization,

$$\bar{w}_p^{(\varepsilon)}(\mathbf{z}) = \frac{w_p(\mathbf{z})}{S(\mathbf{z}) + \varepsilon_{\text{win}}}, \qquad S(\mathbf{z}) = \sum_{q \in \mathcal{P}} w_q(\mathbf{z}),$$

with $\varepsilon_{\text{win}} = 10^{-10}$, except 10^{-12} for the porous topology update. Unlike the exact normalization in Eq. 3, these weights sum to $S/(S + \varepsilon_{\text{win}})$; compact-support zeros remain exact. The full candidate set

remains in the denominator after selection. For a midpoint grid of spacing h in d spatial dimensions, every domain point is within $\sqrt{d}\,h/2$ of a center. The radii of all normalized scaffolds in Table 6 exceed this distance, as do the two smaller learned-solver radii, ensuring $S > 0$ throughout their bounding domains. For the cubic profile, $S \geq [1 - dh^2/(4r_p^2)]^3 > 0$. The unnormalized singular-repair variants have no denominator coverage requirement and contribute zero wherever all active windows vanish.

Coverage of the full candidate scaffold differs from coverage by the selected supports. The selected POU-S supports leave 72%, 75%, and 62% of the L-shaped, crack-tip, and point-source domains unchanged; 34% ($k_R = 2$) and 57% ($k_R = 5, 10$) for the variable-conductivity rows; and 52% for the porous topology update. The reported outside regions of Table 3 lie 84%, 96%, and 85% beyond the selected supports, respectively. For the 100-step FNOs in Table 2, selected supports cover 81.8–100% of the continuous spatial interval for Allen–Cahn/CDR and 100% for Burgers (81.25–100% and 100% of evaluation-grid nodes, respectively). The 350-step comparison covers 100% for Allen–Cahn/Burgers and 98.5–100% for CDR. The historical frozen-solver rows cover the domain (Appendix C.6).

A.4 Packet bases and coefficient counts

Table 7 complements the candidate scaffolds in Table 6 with the local basis and frozen-field features. All reported fields are scalar. Tensor-product rows retain every multi-index with $0 \leq i_j \leq d$ for each of the D packet coordinates, giving $M = (d+1)^D$. Singular-repair rows instead retain $i_1 + i_2 \leq d$, giving $M = (d+1)(d+2)/2$. The packet coefficient count is $K \cdot M \cdot L$, where L is the number of entries in η; listed enrichment and coordination coefficients are additional.

Table 7: **Packet bases and coefficient counts.** The degree d is per coordinate unless marked "total"; M counts polynomial terms, η lists the L conditioning features ($g = g_{\text{in}}$), and K counts active packets. Counts include the stated extras. Characteristic CDR uses $\zeta = x - vt$ in place of x.

Result	Coords.	d	M	η	K	Extra	Count
FNO Allen–Cahn, Burgers, CDR (Tables 2, 23, 24)	x, τ	3	16	$[1, g, g^2]$	6	–	288
FF-MLP POU-D (Table 23)	x, τ	3	16	$[1, g, g^2]$	12	–	576
Singular repair (Tables 3, 26)	x, y	3, total	10	[1]	4 (2, 8)	enrichment: K	40 (44)
Variable conductivity (Table 4)	x, y	5	36	[1]	4	POU-C: 16	144 (160)
Circle/sunflower interfaces (Table 4)	x, y	4	25	[1]	6	POU-C: 16; subdomain: 32	166; 182
Porous topology update (Table 27)	x, y	3	16	[1]	10	–	160
Supervised local defect (Table 29)	x	11	12	$[1, g, g^2]$	2	–	72
1D multiscale (Table 31)	x	3	4	[1], $[1, g]$, $[1, g, g^2]$	8	–	32, 64, 96

For singular repair, the main setting $K = 4$ gives 40 smooth coefficients or 44 with one enrichment per packet; the $K = 2, 4, 8$ budget ablation uses $10K$ or $11K$, respectively. The POU-D rows activate the full scaffold: 576 coefficients for the variable-conductivity and porous benchmarks (16 packets of 36 and 36 packets of 16). The POU-C and subdomain extras are low-order tensor Chebyshev relaxation fields of degree 3 (16 coefficients) on the whole domain or on each side of the interface. The local Fourier variants of Table 31 replace the four Chebyshev terms by $\cos(k\pi\xi)$ and $\sin(k\pi\xi)$, $k \in \{1, 2\}$. The subdomain interface variant uses three packets on each side.

A.5 Packet selection

The main POU-S deployment results select packets before fitting by the windowed residual energy

$$s_p = \sum_{\mathbf{z} \in Q} \left(\bar{w}_p(\mathbf{x})\, r[g_{\text{in}}](\mathbf{z}) \right)^2,$$

where $r[g_{\text{in}}]$ is the residual of the new problem evaluated on the frozen field and Q is the interior collocation set; the K largest scores are activated. The score uses only the PDE, the geometry, and the frozen field. Implementations differ as follows.

- **Frozen FNO and CDR rows.** Q is the interior space–time grid, excluding the spatial boundary and the first and last time levels; the score is averaged over Q, which does not change the ranking. Characteristic-aligned packets evaluate $\bar{w}_p$ in $\zeta = x - vt$. Selection is done once.
- **Frozen FF-MLP rows.** Selection is stagewise: three stages add two packets each, and after each stage's fit the score is recomputed from the residual of the current adapted field, excluding packets already selected. The selection-initialized POU-D fit then optimizes all packets, with stronger ridge on unselected coefficients during initialization (10^{-3} versus 10^{-7}); the subsequent nonlinear polish uses a displacement penalty around that initialization.
- **Singular repair.** The score is normalized by $\sum_{\mathbf{z}\in Q} w_p(\mathbf{z})^2$ and uses the unnormalized Wendland windows; packets are taken in score order, skipping any candidate whose center lies within 0.18 (0.10 for the point source) of an already selected center. The point-source residual includes the discrete source.
- **Interfaces.** Subdomain POU-C selects the three highest-scoring packets on each side of the known interface.
- **Porous topology update.** Q contains the valid solid cells of the new geometry. The topology variant in Table 27 instead uses the known change in geometry, scoring $s_p^{\text{topo}} = \sum_{\mathbf{z}\in B} \bar{w}_p(\mathbf{z})^2$ over the band B of changed solid/void cells dilated by three cells.

The scores are continuous, and no tie-breaking rule is needed. The selection ablations of Table 24 replace s_p as follows: random selection draws K packets uniformly without replacement, once per run; sensor-aware selection uses $\hat{s}_p + 4\,\hat{s}_p^{\text{sens}}$, where hats denote division by the maximum over p and $s_p^{\text{sens}} = \sum_i \bar{w}_p(\mathbf{z}_i)(y_i - g_{\text{in}}(\mathbf{z}_i))^2 / \sum_i \bar{w}_p(\mathbf{z}_i)$ over the sensors; oracle selection uses $\sum_{\mathbf{z}} \bar{w}_p(\mathbf{z})(u^\star - g_{\text{in}})^2(\mathbf{z}) / \sum_{\mathbf{z}} \bar{w}_p(\mathbf{z})$ over the full grid, which requires the dense target and is diagnostic only.

A.6 Adapter fitting details

Algorithm 1 Partition-of-unity adapter fitting

Input	Frozen backbone g, candidate patch family $\mathcal{P}$, local basis $\{\psi_m\}$ for each correction packet, supervised or physics-informed objective $\mathcal{L}$, selection budget K, optional coordination space $\mathcal{R}$, ridge/anchor weights.
Output	Adapted field $u(\mathbf{z}) = g(\mathbf{z}) + A_{\theta_A}(\mathbf{z})$.
1	Evaluate the frozen backbone on the data, collocation, boundary, or initial points required by $\mathcal{L}$.
2	Compute a scoring signal: supervised residual $u_{\text{target}} - g$, physics residual $\mathcal{N}[g]$, boundary/initial mismatch, or a physical-change indicator.
3	Score each patch $p \in \mathcal{P}$ by the windowed energy of this signal (Appendix A.5).
4	Choose the active packet set S: top-K patches for selected POU, all patches for dense POU, or a prescribed refinement layer for stagewise POU.
5	Use the benchmark window convention in Table 5: Eq. 3 retains the full scaffold in the denominator; singular repair instead uses $\bar{w}_p = w_p$ without normalization. Build packet columns $\bar{w}_p(\mathbf{x})\,\eta_\ell(g(\mathbf{z}))\,\psi_m(\boldsymbol{\xi}_p(\mathbf{x}), \tau(t))$ for $p \in S$ and set inactive coefficients to zero.
6	Fit the active coefficients under the global objective $\mathcal{L}(g + A_{\theta_A})$: use ridge or anchored least squares for supervised residuals and linear PDE objectives; use AdamW or a linearized ridge warm start for nonlinear physics objectives.
7	For stagewise variants, repeat steps 2–6 on the current adapted residual and accumulate the selected packet space.
8	Optionally fit the selected/accumulated packets plus $\mathcal{R}$ (POU-C), or activate and fit all packets (POU-D), using the deployment objective and any preservation masks.
9	Return the frozen-backbone adapted field $g + A_{\theta_A}$; the backbone parameters of g remain fixed throughout adapter fitting.

Residual discretization. The FF-MLP runs and calibrated RBF local-damage diagnostic use automatic differentiation of coordinate backbones, with automatic or analytic differentiation of packet factors and product rules for conditioning features. FNO Allen–Cahn, Burgers, and CDR use second-order centered differences on the 32×32 space–time grid, applied to both the frozen field and complete basis columns at interior nodes; boundary/initial conditions enter as trace rows. Singular repair uses the five-point grid Laplacian; porous heat uses a conservative five-point stencil with face-averaged conductivity on valid solid cells. The material-change benchmarks use analytic derivatives of the manufactured frozen field and off-grid centered differences of basis columns (and any boundary lift) in float64: $h = 10^{-4}$ for variable conductivity and $h = 10^{-3}$ for circle/sunflower interfaces, with central normal derivatives and traces evaluated on the respective sides of the interface.

Training and adaptation budgets. Table 2 uses validation-selected FNO fitting budgets, detailed below. Historical experiments retain 180-step FNO and 400-step FF-MLP AdamW baselines with untuned learning rates. Historical FNO POU-S uses one linearized ridge solve; FF-MLP POU-D adds 250 AdamW coefficient steps after ridge initialization. These earlier rows match deployment signals, but not all normalization, discretization, solver, or compute choices. The separate FF-MLP validation sweep tests sensitivity to its fixed FT recipe (Appendix A.10).

A.7 Linearized coefficient fitting

The main method fits a correction in the field family $u = g_{\text{in}} + A_\theta$. For nonlinear PDEs, there are two natural coefficient solvers: optimize the corrected-field residual directly, or solve a first-order defect equation around the frozen prediction. The historical one-shot POU variants use the latter; the primary matched FNO comparison repeats the linearization with a descent safeguard. Let B_S be the matrix of active POU packet columns for a selected support set S, so that $A_\theta = B_S a$, and let $\mathcal{R}$ collect the PDE residual equations at the current collocation points. The one-step linearized proposal solves

$$a_{\text{lin}} = \arg\min_a \|W_{\text{phys}} \left(\mathcal{R}[g_{\text{in}}] + J_{g_{\text{in}}} B_S a\right)\|_2^2 + \|W_{\text{tr}} \left(\mathcal{C}[g_{\text{in}}] + C B_S a - y_{\text{tr}}\right)\|_2^2 + \lambda_{\text{ridge}} \|a\|_2^2, \tag{16}$$

where $J_{g_{\text{in}}}$ is the residual Jacobian evaluated at the frozen prediction, and $\mathcal{C}$ stacks boundary, initial, sparse-sensor, interface, or preservation-anchor quantities with prescribed values y_{tr}. The matrix C is the Jacobian of $\mathcal{C}$ at g_{in}; for linear trace constraints, it is the trace operator itself. Among the historical one-shot fits, singular repair scales the proposal by the largest $\alpha \in \{1, 0.75, 0.5, 0.375, 0.25, 0.1875, 0.125, 0.0625, 0.03125\}$ satisfying the preservation and benchmark-specific acceptance tests in Appendix B.2; if none passes, the frozen field is kept. The guarded region includes low-weight window tails and is not the strict support exterior. The same test is applied to the global control in the crack-tip and point-source benchmarks. These tests use only the frozen field, the PDE, and the sensors. Other historical one-shot rows apply the solved coefficients without damping; the FF-MLP POU-D rows then refine them with 250 AdamW steps. The matched FNO GN comparison instead uses the Armijo safeguard specified in Appendix A.9. This choice changes only the coefficient solve inside POU-S, POU-C, or POU-D; it does not introduce a fourth operating mode.

For linear PDE residuals, Eq. 16 reduces to the same least-squares residual correction obtained from the corrected-field objective. For nonlinear residuals, this first-order approximation need not produce the same update as direct corrected-field optimization. This subsection specifies the fitting formulation used by the least-squares variants; it does not report a comparative solver ablation.

A.8 Concrete relaxation solves

The coordinated variants in the experiments keep the backbone fixed. Relaxation means adding a small correction space to the selected local packets and solving the enlarged adapter problem jointly; it does not fine-tune g. Let $A_S(\mathbf{z}) = B_S(\mathbf{z}) a_S$ denote the selected POU correction built from the active patch set S, and let $R(\mathbf{z}) = B_R(\mathbf{z}) a_R$ denote the full-domain relaxation field, with basis B_R and fitted coefficients a_R. The relaxed field is

$$u_{\text{relax}}(\mathbf{z}) = g(\mathbf{z}) + A_S(\mathbf{z}) + R(\mathbf{z}). \tag{17}$$

This adapter-space relaxation is distinct from backbone relaxation. In separate baselines and hybrid variants, we also compare against full fine-tuning, last-layer tuning, LoRA, or LoRA+POU when

the backbone architecture supports it. Those variants update parameters of g or add weight-space adapters, and are used to mark the scope boundary: POU is intended as a compact, spatially inspectable correction interface, while unrestricted backbone relaxation is a stronger but less localized adaptation mechanism.

For a linear adapter basis $B = [B_S \; B_R]$, the calibrated Poisson relaxation diagnostic in Appendix D.2 uses physics, boundary, and ridge terms plus a field-magnitude penalty on the correction. It uses the same least-squares system as selected POU, but with the concatenated basis:

$$\min_a \|\mathcal{N}[Ba] + r_{\text{base}}\|_2^2 + \lambda_B \, \|\mathcal{B}[Ba] + b_{\text{base}}\|_2^2 + \lambda_A \|Ba\|_2^2 + \lambda_{\text{ridge}} \|a\|_2^2. \tag{18}$$

Here $a = (a_S, a_R)$, and r_{base} and b_{base} are the calibrated residual and boundary mismatch defined in Appendix D.2. The field penalty penalizes large correction fields over the domain. The supervised guidance sweep in Appendix D.3 instead fits the correction to dense target values.

A.9 Pretraining, solver settings, and validation

The principal FNO experiment uses 32 training, eight validation, and 12 test instances per PDE. For Burgers and Allen–Cahn, split seed 20260925 plus a PDE offset stratifies the datasets by viscosity and reaction strength, respectively. CDR uses 52 manufactured parameter tuples generated with seed 20260926, shuffled within four spatial-mode strata to assign eight/two/three instances per mode to training/validation/test. Saved manifests fix all parameter tuples and IDs; test and validation instances are excluded from backbone training. Seeds 11, 23, and 37 are reused across the same test instances. The 32-instance training regime is a controlled small-data setting. Table 2 uses the 100-step checkpoints; Table 14 retains the 350-step comparison from the same trajectories and cohort. Adaptation recipes are selected separately at each strength. The CDR 100-step checkpoints are recovered by replaying the original training trajectory and checking its 350-step predictions against the archive. Both comparisons use default spatial packets, independently of the historical characteristic-coordinate study.

Table 8: **Pretraining and evaluation units.** Supervised pretraining precedes deployment adaptation. Dense validation fields select offline optimization recipes; dense test fields are withheld from adaptation.

Backbone	Architecture	Pretraining	Evaluation
FNO, matched comparison	width 20, depth 3, 10 modes; 32×32 grid	32 instances; 100/350 AdamW steps, learning rate 2×10^{-3}, decay 10^{-6}	8 validation and 12 test instances per PDE; seeds 11, 23, 37
FF-MLP	12 Fourier frequencies; two width-48 hidden layers	Per-instance 28×28 samples; 800 supervised steps, then 200 physics/trace steps with dense anchor weight 20	Same-instance refinement; 12 instances, 3 seeds; 64×64 evaluation
Historical FNO, CDR	width 20, depth 3, 10 modes; 32×32 grid	32 generated training instances; 350 AdamW steps, learning rate 2×10^{-3}	12 distinct manufactured test instances; seeds 11, 17, 23; fixed-budget adaptation

FF-MLP pretraining uses AdamW with learning rates 10^{-3} for the supervised stage and 5×10^{-4} for the physics/anchor stage, with weight decay 10^{-6} in both. Each validation or test instance has its own imperfect pretrained field; this study measures same-instance refinement rather than operator generalization.

Shared FNO objective and coordinate controls. For Table 2, all adapted methods minimize

$$J(u) = \text{MSE}_{\text{PDE}} + 25\,\text{MSE}_{\text{IC}} + 4\,\text{MSE}_{\text{BC}} + 25\,\text{MSE}_{\text{sens}} + 10^{-6}\,\text{mean}(u-g)^2,$$

omitting the sensor term when $m = 0$. Boundary MSE averages the two sides. The global control uses tensor Chebyshev degrees 23 in space and 3 in time with features $[1, g, g^2]$, giving 288 coefficients; POU uses six degree-3 packets with the same features and coefficient count. Columns are RMS-scaled and transformed to a field-orthonormal basis by SVD, retaining singular values above 10^{-10} times the

largest. The 288-coefficient budget is nominal: retained POU ranks in Table 2 are 278–288 for Allen–Cahn, 286–288 for Burgers, and 279–288 for CDR; the global rank is 288 throughout. The shared penalty acts on the physical correction field, not differently scaled raw coefficients. Both coordinate spaces use augmented least squares (relative cutoff 10^{-12}) within safeguarded Gauss–Newton: at most 12 linearizations, Armijo constant 10^{-4}, up to 16 halvings, gradient tolerance 10^{-8} and relative objective-change tolerance 10^{-9}. This comparison changes feature conditioning, numerical scaling, and fitting jointly relative to the historical pipeline; it is not an isolated normalization ablation.

Validation-only budget selection. Global and POU budgets are chosen from $\{1, 2, 4, 8, 12\}$ linearizations. Full FT uses Adam without an additional parameter-weight penalty, learning rates $\{10^{-4}, 10^{-3}\}$, and checkpoints $\{50, 100, 180, 350, 700, 1400, 2800, 5600\}$. For each PDE/sensor condition, the recipe minimizes mean relative error on eight validation instances and three seeds; ties favor fewer steps, then lower learning rate. FT selection considers positive step counts; the frozen column reports the separate zero-step option. For the 100-step backbones, both learning rates run through 5600 validation steps in every condition. Selected FT recipes are 10^{-4}/50 and 10^{-4}/5600 for Allen–Cahn at $m = 0$ and 32, 10^{-3}/2800 for both Burgers conditions, and 10^{-4}/5600 and 10^{-4}/2800 for CDR. Global/POU budgets are 2/2 and 4/2 for Allen–Cahn, 8/8 for Burgers, and 1/1 for CDR. The sensor-free Allen–Cahn frozen validation error, 0.09314, is below the best positive FT value, 0.11337; allowing zero steps would select the frozen field. These bounded searches do not establish convergence.

The earlier 350-step comparison in Table 14 used a validation rule stopping after two checkpoints without a 1% improvement, no earlier than 700 steps. For sensor-free Allen–Cahn at 10^{-4}, this stopped a recovering trajectory at 700 and selected 50 steps. A validation-only continuation to 5600 reproduced earlier predictions exactly and reduced validation error from 0.02008 to 0.01477, below the frozen value 0.01547; selection was relocked before test replay. This differs from the 100-step result above. The earlier cap had been reduced from 11,200 to 5,600 after validation and a partial test run for resource reasons, before relocking selection and rerunning all 144 Allen–Cahn/Burgers test cells. Test outcomes did not enter recipe selection. The 350-step FT recipes use 10^{-4} except Burgers with sensors (10^{-3}); both CDR learning rates run through 5600. The two backbone strengths and preceding robustness/control studies reuse the same cohort and are not independent replications.

Pipeline-specific scaling. For a residual block $A_b a - d_b$ with n_b rows, block-mean weighting scales both sides by $\sqrt{w_b/n_b}$; summed weighting uses $\sqrt{w_b}$. These conventions are not interchangeable. Table 9 specifies the other normalizations and regularizers. Their weights must be interpreted together with their block scaling.

The enriched corner row replaces the strong-form PDE block by a gradient-energy block, $\nabla(g + Ba)$, with the same weight 0.1 and boundary/sensor/trust terms; its guard checks energy. Its raw PDE residual is therefore an evaluation diagnostic, not the fitted energy. The material/interface normalization depends on the basis matrix; it does not yield the common field objective used in the matched FNO experiment. Degenerate blocks with $s_b \leq 10^{-12}$ are left unscaled. The FF-MLP ridge initializes the all-packet fit; its subsequent polish penalizes mean squared displacement from that initialization, rather than retaining the diagonal selected/unselected ridge in the final objective. Thus the archived FF-MLP "POU-C" is labelled POU-D (selection-initialized) here. Its archived "POU-D" counterpart uses uniform 10^{-7} ridge initialization in the same 576-dimensional packet span. FF-MLP FT uses a 28×28 grid while POU fitting uses the 64×64 evaluation grid; their deployment signals agree, but their discretizations and regularizers are not identical. The matched FNO experiment provides the shared-objective comparison.

Timing contract and provenance. Table 2 measures repeated use on an AMD EPYC 7742 / NVIDIA RTX 4090 (24GB) cluster host, with one allocated GPU per PDE. Table 14 and the separate first-use profiles (Table 13) use the earlier AMD Ryzen Threadripper 2950X / NVIDIA RTX 2070 SUPER (8GB) workstation. Both use single-thread CPU coordinate solves and GPU FNO/FT in float64; times are compared within a hardware setup. Pretrained weights and compiled shape-compatible kernels are resident. Each trial includes input/output transfers, the frozen forward pass, packet selection, basis construction, SVD and derivative assembly or optimizer initialization, fitting, and final prediction. Basis factorizations are recomputed for every trial; GPU work is synchronized. We use three repetitions in fixed shuffled method order, report the median per case, then average

Table 9: **Solver settings for the principal and historical comparisons.** PDE, initial, boundary and sensor weights are ordered as applicable. Numerical settings are fixed within each experiment.

Pipeline	Block scaling and weights	Regularization and fitting
Matched FNO (Table 2)	Block means; PDE/IC/BC/sensor $1/25/4/25$	Field penalty 10^{-6}; SVD-scaled safeguarded GN or validation-selected Adam FT, as above
Historical FNO/CDR	Coordinate fits: weighted sums; FT: block means; $1/25/4/25$ (no sensors for CDR)	Coordinate ridge 10^{-6}, one linearized normal-equation solve; FT 180 AdamW steps, LR 10^{-3}, decay 10^{-6}
FF-MLP	Ridge initialization: weighted sums; polish/FT: block means, PDE/IC/BC $1/25/4$	Normal-equation ridge 10^{-7} on selected packet coefficients and 10^{-3} elsewhere; 250-step AdamW polish, LR 2×10^{-3}, mean coefficient displacement penalty 10^{-5}
Singular repair	Each block scaled by $w_b/\max(\mathrm{RMS}(d_b),1)$, so weights enter squared	Augmented least squares, ridge 10^{-7} and the deployment guard in Appendix A.7; sensor weight 100, PDE weight 0.1; boundary/trust 12/0.55 (corner), 20/0.5 (crack/source); source-aware annular/flux weights 0.5/25
Bulk-only κ diagnostic	Each block divided by $s_b=\sqrt{n_b^{-1}\sum_i\|(A_b)_{i,:}\|_2^2}$, then scaled by $\sqrt{w_b}$; PDE/BC/field-trust $1/50/0.01$	Normal-equation ridge 10^{-7}; interface value and flux multipliers both zero
Circle/sunflower interfaces	Same matrix-row normalization; PDE/BC/value/flux/field-trust $1/50/1/1/0.001$	Normal-equation ridge 10^{-7}; joint or subdomain-coupled fitting; both interface blocks active

across cases/seeds. Disk loading, serialization, offline error scoring, and intermediate curve probes are excluded. These are implementation-specific CPU/GPU costs, not inference latencies.

The historical FNO Allen–Cahn/Burgers values 21/25 ms use archived fitting timers, while CDR 61 ms also includes selection and basis construction; both exclude common frozen-field evaluation. Historical FF-MLP timers include the final ridge assembly and coefficient polish but omit common frozen derivatives and prerequisite stage-selection fits. These figures are retained only with their original scope in Table 23; no end-to-end speed factor is inferred from them. Complete FF-MLP replays are reported separately with the validation-selected baseline.

The singular and material/interface costs retain their fit-stage scope. The κ times are medians of three single-thread repeats of the float64-corrected operator under the same fit-stage timer; common field construction and diagnostics are excluded.

Paired uncertainty. Table 11 reports geometric mean paired error ratios with 95% percentile intervals from 5,000 crossed bootstrap resamples of the three backbone seeds and 12 instance IDs. Ratios below one favor POU. Table 12 exposes variation across the three backbone seeds. These intervals are exploratory with only three seeds; grid points, sensor counts, and checkpoints do not increase the number of independently trained backbones.

A.10 FF-MLP VALIDATION AND TIMING PROTOCOL

We retain the original 12 test instances and seeds 11, 23, 37 for each PDE and select an FT recipe on eight separate metadata-stratified validation instances. Validation sampling uses seed 20260930 plus a PDE offset and excludes the FF-MLP test IDs and existing FNO validation/test IDs. Each instance has its own frozen FF-MLP, pretrained by the procedure in Table 8. The sweep crosses learning rates $\{5\times10^{-5},5\times10^{-4},5\times10^{-3}\}$ with checkpoints $\{100,400,800,1600\}$, using all three seeds. One recipe per PDE minimizes mean validation relative error; ties favor fewer steps, then lower rate. All 576 validation predictions and 72 test case/seed pairs are retained. Dense fields serve offline validation selection and evaluation; adaptation uses PDE and initial/boundary traces. This is a bounded sensitivity study, not a claim of globally optimal or converged fine-tuning.

Table 10: **FF-MLP recipes selected before the test replay.** Each validation mean uses eight instances and three initialization seeds.

PDE	FT learning rate	Steps	Validation rel-L_2
Allen–Cahn	5.00×10^{-5}	100	0.207
Burgers	0.0005	800	0.466

The POU recipe is unchanged: three two-packet selection stages initialize a dense 576-coefficient fit over all 12 packets, followed by 250 AdamW coefficient steps. Thus its operating mode is POU-D despite the archived POU-C name. Both methods use automatic derivatives. FT fits on a 28×28 grid and POU on a 64×64 grid. Evaluation uses 64×64 points; the PDE mask excludes $t = 0$ and, for Allen–Cahn, spatial boundaries. Burgers includes the spatial trace points. No sensor or interface terms apply.

Timing uses the workstation in Appendix A.9, with CPU basis assembly/linear solves and GPU neural derivatives and optimization. Both methods reuse shape-compatible compiled kernels. Every POU trial rebuilds selection, basis derivatives and ridge systems; derivatives of the $[1, g, g^2]$ columns use exact product rules with the same autodifferentiated frozen field. FT includes a frozen forward pass, fitting-buffer transfers, optimizer initialization, training and final prediction; POU additionally includes all prerequisite selection fits and the final coefficient polish. Pretraining, compilation warmup, disk operations and offline scoring are excluded. The optimized derivative and polish implementation was checked against the original routine and the archived errors across every test case; validation recipes were kept fixed. The initial replay with per-call POU compilation is preserved separately and is not mixed into these warm timings.

Test errors, paired comparisons, and the fixed-budget sensitivity summary appear in Appendix C.2.

B Main benchmark specifications

The benchmark families follow Sections 4.1, 4.2, and 4.3: frozen-model adaptation, localized/singular errors, and material changes with nonlocal responses. The subsections give each family's equations and setup; Appendix A.1 defines the evaluation metrics. The experiments use manufactured or simulation-derived fields to isolate correction behavior rather than introduce a new dataset benchmark.

B.1 Adapting frozen learned solvers

The imperfect-solver application tests whether POU adapters can improve a frozen prediction on a new instance using residual, initial/boundary traces, and optional sparse measurements, without dense interior target fields for fitting.

Fourier-feature MLP setup. These diagnostics use supervised-pretrained Fourier-feature MLP backbones that are frozen before adaptation. During adaptation, full-network PINN fine-tuning, global controls, and POU adapters are trained from physics residuals with boundary or initial-condition terms; no interior supervised values are used for adapter fitting. These experiments test whether PDE residuals can drive localized POU correction when the physics objective identifies the desired solution branch.

Viscous Burgers. For $(x, t) \in [-1, 1] \times [0, T]$, use

$$u_t + uu_x - \nu u_{xx} = 0, \qquad u(x, 0) = u_0(x), \qquad u(\pm1, t) = b_\pm(t). \tag{19}$$

This benchmark tests nonlinear advection and shock-like residual localization. The default fitting signal is PDE residual plus IC/BC traces; optional sparse measurements are used only in explicitly labeled sparse settings.

Allen–Cahn. For $(x, t) \in [-1, 1] \times [0, T]$, use

$$u_t - \epsilon^2 u_{xx} + \lambda(u^3 - u) = 0, \qquad u(x, 0) = u_0(x), \qquad \mathcal{B}[u] = b. \tag{20}$$

This benchmark tests whether residual-selected local packets remain useful when the error comes from reaction–diffusion dynamics and phase-interface motion rather than pure transport.

Convection–diffusion–reaction. For $(x,t) \in \Omega \times [0,T]$, use

$$u_t + \mathbf{v} \cdot \nabla u - \nabla \cdot (D \nabla u) + ru = f, \qquad u(\mathbf{x},0) = u_0(\mathbf{x}), \qquad \mathcal{B}[u] = b. \tag{21}$$

This benchmark tests linear transport–diffusion physics. The matched comparison uses spatial packets; the historical geometry study compares these with characteristic-aligned packets.

Instances and reference fields. The Burgers and Allen–Cahn instances come from precomputed datasets on $[-1,1] \times [0,1]$ ($T = 1$), stored on a 64×64 grid. Burgers uses $u_0(x) = A \sin(k\pi x)$ with $A \in [0.5, 2.3]$, $k \in [0.5, 2.5]$, and $\nu \in [0.01, 0.3]$ drawn uniformly (143 instances), integrated with SciPy's `solve_ivp` on a 512×512 grid; since k need not be an integer, the boundary values are not zero, and $b_{\pm}(t)$ are taken from the reference solution. Allen–Cahn uses $u_0(x) = (1 - x^2)\cos(\pi x)$, $u(\pm 1, t) = 0$, $\epsilon \in [1.1 \times 10^{-4}, 9.8 \times 10^{-4}]$, and $\lambda \in [1.0, 5.0]$ (200 instances). The CDR instances use $\Omega = [-1,1]$, $T = 1$, $f = 0$, and the manufactured solution $u = e^{-Rt}\left[1 + a\, e^{-D\omega^2 t} \sin(\omega(x - vt))\right]$ with $\omega = m\pi$, from which the initial and Dirichlet data are taken; the 12 historical geometry-study tuples have $v \in [0.7, 2.5]$, $D \in [0.005, 0.018]$, $R \in [0.2, 1.05]$, $a \in [0.35, 0.6]$, and $m \in \{1, \ldots, 4\}$. The matched CDR study uses a separate 52-tuple pool and the stratified split in Appendix A.9; its 12 test tuples span $v \in [0.699, 2.193]$, $D \in [0.00858, 0.01745]$, $R \in [0.328, 0.945]$, $a \in [0.368, 0.582]$, and all four modes. Its 32 sensors are sampled uniformly without replacement from interior residual-grid nodes, with seed $20260926 + 1009i$ for case index i, and shared across backbone seeds. The FNO rows are evaluated on a 32×32 grid (bilinear resampling of the dataset fields; the CDR formula is evaluated directly), and the FF-MLP rows on the 64×64 dataset grid.

B.2 Localized and singular errors

The localized/singular-error application starts from a coarse numerical field $g_{\mathrm{dam}} = \mathcal{I}_H^h u_H$ that is accurate away from a localized singular region but underresolved near that region. The key metrics are local rel-L2 on the singular neighborhood D_{loc}, global rel-L2, and outside-region rel-L2 on a separate evaluation mask D_{out}. The fine-grid reference is used only after adaptation for evaluation, except for oracle floor diagnostics explicitly marked as dense-target fits.

For the crack-tip and point-source problems, D_{loc} and D_{out} are separated radial masks centered at the singularity: $r \le 0.36$ and $r \ge 0.55$ for the crack tip, and $r \le 0.16$ and $r \ge 0.30$ for the point source. The intervening annuli, $0.36 < r < 0.55$ and $0.16 < r < 0.30$, are excluded from both regional metrics but included in the full-domain metric reported in Table 26. Crack-tip regional masks retain outer-boundary and crack-trace grid points; its global metric excludes the crack trace. All point-source error masks exclude outer-boundary points. Reported errors are averaged over four coarse-grid resolutions and three sensor-placement seeds.

L-shaped Poisson corner. The first benchmark uses the classical re-entrant L-shaped domain $\Omega_L = (-1,1)^2 \setminus \{(x,y) : x > 0,\ y < 0\}$. The reference is the harmonic corner singularity

$$u^\star(r,\theta) = r^{2/3} \sin(2\theta/3), \qquad -\Delta u^\star = 0 \quad \text{in } \Omega_L, \qquad u = u^\star \quad \text{on } \partial\Omega_L, \tag{22}$$

with polar coordinates centered at the re-entrant corner. The coarse field is a finite-difference solve with exact Dirichlet boundary data, interpolated to the fine grid. On the L-shaped evaluation grid, including its boundary, $D_{\mathrm{loc}} = \{r \le 0.48\}$ and $D_{\mathrm{out}} = \{\max_{p \in \mathcal{S}} w_p \le 0.08\}$, where $\mathcal{S}$ is the selected packet set and w_p are the unnormalized Wendland windows. The same masks are used for Frozen, Global, POU-S, and the singular-enriched variant within each case. The outside mask includes low-weight window tails, so it is not the strict support exterior and its error can change after local adaptation.

Crack-tip Poisson. The second benchmark uses a square slit domain $\Omega_C = (-1,1)^2 \setminus \{(x,0) : 0 \le x < 1\}$. The reference is the crack-tip harmonic field

$$u^\star(r,\theta) = r^{1/2} \sin(\theta/2), \qquad -\Delta u^\star = 0 \quad \text{in } \Omega_C, \qquad u = u^\star \quad \text{on the outer and crack traces.} \tag{23}$$

The coarse solver underresolves the singular gradient at the crack tip while preserving the far field. The problem-aware variant augments the smooth POU basis with the crack-tip singular shape inside the selected packets.

Point-source Poisson. The third benchmark uses a source-driven singularity on the unit square,

$$-\Delta u = \delta_{\mathbf{x}_0} \quad \text{in } \Omega, \qquad u = 0 \quad \text{on } \partial\Omega, \qquad \Omega = (0,1)^2, \quad \mathbf{x}_0 = (0.5, 0.5). \tag{24}$$

The reference is the fine-grid discrete Green's function, and the input field is the corresponding coarse-grid discrete Green's function interpolated to the fine grid. The singular source node is excluded from pointwise residual fitting. The problem-aware variant uses Green/log-enriched packets, annular residual equations around the source, and the weak Gauss-law constraint

$$-\int_{\partial B_\rho(\mathbf{x}_0)} \nabla u \cdot n \, ds = 1, \qquad \rho = 0.12. \tag{25}$$

This weak constraint encodes the source strength without forcing the adapter to chase a grid-dependent pointwise residual at the delta node.

Grids, solvers, and enrichment. Coarse fields are computed on $n_H \times n_H$ grids with $n_H \in \{7, 9, 11, 13\}$ by the five-point finite-difference Laplacian and transferred to the fine grid (129×129 for the L-shaped domain, 161×161 otherwise) by piecewise-linear interpolation, after which the exact boundary values are restored. Polar angles are taken in $[0, 2\pi)$, so the branch cut lies along the positive x-axis: on a domain edge for the L-shaped corner and on the slit for the crack, whose two faces share one row of Dirichlet nodes with exact value 0. The L-shaped and crack-tip references are the analytic fields; the point-source reference is the fine-grid discrete Green's function with $f = 1/h^2$ at the node nearest $\mathbf{x}_0$, and residual rows within $1.5h$ of the source are excluded from fitting. The problem-aware variants add one enrichment column $w_p\phi$ per selected packet, with $\phi = r^{2/3}\sin(2\theta/3)$ (L-shaped), $\phi = r^{1/2}\sin(\theta/2)$ (crack), and, for the point source, $\phi = -\log(\max(r, h))/(2\pi)$, set to zero on $\partial\Omega$ and scaled to unit maximum. The L-shaped variant also uses a gradient-energy objective. The point-source variant adds residual rows on the annulus $0.035 \le r \le 0.18$ and the flux condition of Eq. 25, evaluated at 256 equally spaced points with central differences of width 0.03 in the normal direction. These variants therefore use the singular location, the singular exponent and angular form, and the unit source strength as problem information.

Preservation priors and acceptance. Let $W = \max_{p \in S} w_p$ and $U = \{W > 0.08\}$, using the selected unnormalized windows. Fitting penalizes the correction on interior residual-valid nodes outside U whose absolute frozen residual is at most the q-quantile over the full interior residual grid, before restricting the PDE-fitting rows, with $q = 0.55$ for the corner and $q = 0.50$ for crack/source. These low-residual masks express a preservation prior, not known-correct target values.

The acceptance guard measures $\mathrm{RMSE}_G(u - g)$ on G, the domain grid outside U for the corner, the non-crack grid outside U for the crack, and interior points with $r \ge 0.30$ outside U for the source. It requires

$$\mathrm{RMSE}_G(u-g) \le \begin{cases} \max\{0.035,\ 0.055\max(\mathrm{RMSE}_G(g), 1)\}, & \text{corner/crack,} \\ \max\{0.002,\ 0.08\,\mathrm{RMSE}_G(g)\}, & \text{source.} \end{cases}$$

The corner additionally requires active-region PDE RMSE at most 0.995 of the frozen value for residual fits, or gradient-energy RMSE at most 0.999 for energy fits; the active region is the residual-valid part of U, falling back to all residual-valid nodes if fewer than 20 remain. The source requires nonincreasing sensor RMSE when sensors are present; the crack has no additional residual test. Crack-tip and point-source fits restore boundary traces exactly; corner fits impose weighted boundary constraints. The damping schedule is given in Appendix A.7. These guard and fitting masks can include nonzero window tails; they serve different purposes from the evaluation mask D_{out} and the strict support exterior, although some masks coincide in the corner benchmark.

B.3 Structured changes with nonlocal response

The structured, nonlocal-response application uses manufactured fields under heterogeneous operators. The variable-κ rows diagnose bulk-residual correction from a prescribed smooth field; the circle/sunflower experiments adapt an old material-instance solution with explicit interface constraints.

Variable-conductivity Poisson/heat. For a two-material domain $\Omega = \Omega_1 \cup \Omega_2$ with interface Γ, use

$$-\nabla\cdot(\kappa_{\text{new}}(\mathbf{x})\nabla u) = f \quad \text{in } \Omega_1 \cup \Omega_2, \qquad \mathcal{B}[u] = b \quad \text{on } \partial\Omega. \tag{26}$$

The physical interface conditions are

$$[\![u]\!]_\Gamma = 0, \qquad [\![\kappa_{\text{new}}\nabla u \cdot n]\!]_\Gamma = 0. \tag{27}$$

For the reported κ diagnostic, $\Omega = [0,1]^2$, $\Gamma = \{x = 1/2\}$, $\kappa_{\text{left}} = 1$ and $\kappa_{\text{right}} = k_R \in \{2, 5, 10\}$. The frozen field is $g = (10/7)x(1-x)(1+0.8x)\sin(\pi y)$. The manufactured target is $u^\star = \phi(x)\sin(\pi y)$, where $\phi = x$ on the left and $\phi = (2 + 1/k_R)(1-x) - (2 + 2/k_R)(1-x)^2$ on the right; forcing is $f = -\kappa(\phi'' - \pi^2\phi)\sin(\pi y)$ within each material. This prescribes the new operator and forcing and does not establish a conductivity-only change with fixed forcing. The 48×48 bulk grid excludes boundaries and a band $|x - 1/2| \le 2h$, $h = 10^{-4}$. Both interface multipliers are zero in Table 4; hence Eq. 27 describes the target transmission conditions, not constraints imposed by those fits. Interface diagnostics use 48 trace locations at $x = 1/2 \pm 5 \times 10^{-4}$ and inherit finite-offset error.

Paper-formula elliptic interface benchmarks. Following the hard-constraint elliptic-interface benchmarks of He et al. (2026), we also use manufactured two-material problems on $\Omega = [-1,1]^2$ with an interface Γ separating an inner and outer material. The residual is

$$-\nabla\cdot(\alpha(\mathbf{x})\nabla u) = f \quad \text{in } \Omega \setminus \Gamma, \qquad u = b \quad \text{on } \partial\Omega, \tag{28}$$

with prescribed interface jumps

$$[\![u]\!]_\Gamma = q_u, \qquad [\![\alpha\nabla u \cdot n]\!]_\Gamma = q_\alpha. \tag{29}$$

The circular case uses $\Gamma = \{r = r_0\}$ with $r_0 = 0.5$ and the piecewise reference

$$u_{\text{in}}(r) = \frac{r^3}{\alpha_{\text{in}}}, \qquad u_{\text{out}}(r) = \frac{r^3}{\alpha_{\text{out}}} + \left(\frac{1}{\alpha_{\text{in}}} - \frac{1}{\alpha_{\text{out}}}\right) r_0^3. \tag{30}$$

The sunflower case uses $r_\Gamma(\theta) = 0.5 + \frac{1}{7}\sin(5\theta)$ and

$$u_{\text{in}}(r) = \frac{r^2}{\alpha_{\text{in}}}, \qquad u_{\text{out}}(r) = \frac{r^4 + 0.1\log(2r)}{\alpha_{\text{out}}}. \tag{31}$$

For the circle/sunflower material-swap experiments, the frozen field is the old manufactured solution and the new problem swaps the material contrast. The inside/outside contrasts $(\alpha_{\text{in}}, \alpha_{\text{out}})$ change from $(1, 1000)$ to $(1000, 1)$ for the circle and from $(1, 10)$ to $(10, 1)$ for the sunflower. Both use a 97×97 grid and interface traces offset by 0.00729167 on each side, distinct from the 10^{-3} derivative step. If the swap changes the outer boundary trace, the boundary-lifted procedure first adds an analytic square-boundary lift matching the new minus old Dirichlet trace on $\partial\Omega$ (the Coons-type transfinite interpolation of the trace difference from the four edges). The Frozen entry in Table 4 is the old solution plus this lift, every other method starts from the same lifted field, and exact preservation outside the active supports is relative to it; POU coefficients are then fit from the new residual and the interface value/flux constraints. Subdomain POU-C keeps separate inside/outside packet sets and couples them through the same interface terms. The reported circular subdomain row uses six Schwarz-style alternating sweeps; the sunflower row uses a joint solve.

B.3.1 Porous topology-update settings

Porous topology update. This benchmark uses paired two-dimensional steady heat solutions where the outer body and boundary conditions are held fixed while the internal void layout changes. Solution A serves as the frozen anchor field g_{old}. For porous heat on configuration B, we evaluate

$$-\nabla\cdot(\kappa_{\text{new}}\nabla u) = f \quad \text{on valid solid cells}, \qquad \mathcal{B}_{\text{outer}}[u] = b, \qquad \mathcal{B}_{\text{void}}[u] = b_{\text{void}} \tag{32}$$

with the residual masked away from invalid void cells. In the deployment setting, the adapter is fit from the new geometry and topology mask, conductivity map, finite-difference residual on valid solid cells, outer-boundary traces, optional sparse measurements, and a weak preservation prior (weight 0.2, on every second cell) that keeps the correction small on solid cells shared by both geometries and at least three cells from the changed region. This mask is determined by the geometry alone and acts as a regularizer, not as known correct values; the dense B solution is reserved for evaluation. This benchmark evaluates topology-induced scope limits where local geometric modifications require sparse anchoring or coordinated capacity to capture nonlocal thermal responses. Additional topology-update results and representative heat fields appear in Appendix C.8.

C SUPPORTING DEPLOYMENT RESULTS AND ABLATIONS

The supporting evidence begins with the primary matched FNO comparison and validation-selected FF-MLP refinement, followed by solution and constraint diagnostics, historical comparisons, and ablations. Pretraining, solver choices, validation selection, and timing procedures are collected in Appendix A; benchmark equations and settings are in Appendix B.

C.1 MATCHED FNO COMPARISON

These 100-step-backbone results support Table 2. Pretraining, validation selection, uncertainty estimation, and timing procedures are specified in Appendix A.9. Table 11 reports paired relative-error ratios (values below one favor POU), and Table 12 reports means for each pretrained backbone seed.

Table 11: **Paired relative-error comparisons for the 100-step backbones.** All test cases are retained. "Wins" counts lower POU error within each case/seed pair.

PDE	m	Comparator	POU/comparator	95% interval	POU wins
Allen–Cahn	0	Global	2.023	$[1.158, 3.343]$	10/36
Allen–Cahn	0	FT	0.249	$[0.121, 0.482]$	35/36
Allen–Cahn	0	Frozen	0.323	$[0.158, 0.611]$	31/36
Allen–Cahn	32	Global	1.312	$[0.964, 1.982]$	15/36
Allen–Cahn	32	FT	0.461	$[0.278, 0.835]$	34/36
Allen–Cahn	32	Frozen	0.135	$[0.084, 0.213]$	36/36
Burgers	0	Global	1.162	$[1.079, 1.260]$	3/36
Burgers	0	FT	1.624	$[1.328, 2.089]$	0/36
Burgers	0	Frozen	0.120	$[0.083, 0.168]$	36/36
Burgers	32	Global	1.150	$[1.087, 1.228]$	2/36
Burgers	32	FT	1.589	$[1.318, 1.915]$	1/36
Burgers	32	Frozen	0.104	$[0.072, 0.145]$	36/36
CDR	0	Global	1.186	$[1.009, 1.819]$	14/36
CDR	0	FT	6.049	$[3.461, 10.594]$	1/36
CDR	0	Frozen	0.380	$[0.176, 0.658]$	36/36
CDR	32	Global	1.231	$[1.011, 2.090]$	16/36
CDR	32	FT	6.556	$[3.250, 11.652]$	3/36
CDR	32	Frozen	0.288	$[0.115, 0.567]$	36/36

Accuracy versus time. Figure 3 retains every tested budget. Marked endpoints are selected on validation data; the test curves are descriptive and are not used to select an individual test case's budget.

Sensitivity to backbone pretraining. Table 14 retains the 350-step results on the same 12 test instances, three backbone seeds, and sensor sets as the 100-step comparison in Table 2. Each strength selects its own adaptation recipe on validation data. Reducing pretraining to 100 steps increases mean frozen error by $6.69\times$ for Allen–Cahn and about $2.45\times$ for Burgers/CDR. POU improves the weaker frozen predictions in 211 of 216 case/seed/sensor pairs, but its final mean error is higher than with 350-step backbones in every condition. Global Cheb. is more accurate than POU in all six 100-step conditions. Thus the correction remains useful as starting accuracy falls, with performance depending on the backbone and problem. The two hardware setups preclude attributing differences in measured time to pretraining strength alone.

Fixed-recipe single-instance diagnostic. The earlier diagnostic below compares the same saved 100-step and 350-step FNO checkpoints on one fixed Allen–Cahn validation instance (case 74), using seeds 11, 23, and 37. The instance was fixed before inspecting its undertrained predictions. Observations and fitting settings are shared: Global Cheb. and POU-S use 288 nominal coefficients and two GN steps; FT uses learning rate 10^{-4} and 5600 steps. These budgets are not retuned for the earlier checkpoint. Frozen-field features and packet selection are recomputed for each checkpoint; dense interior targets are used only for error evaluation.

Table 12: **Per-seed mean errors for the 100-step backbones.** Each entry averages the same 12 test instances.

PDE	m	Method	Seed 11	Seed 23	Seed 37
Allen–Cahn	0	Frozen	0.0825	0.0952	0.0926
Allen–Cahn	0	Global	0.0202	0.0247	0.0215
Allen–Cahn	0	POU-S (GN)	0.0336	0.0566	0.0672
Allen–Cahn	0	FT	0.1011	0.1207	0.1226
Allen–Cahn	32	Frozen	0.0825	0.0952	0.0926
Allen–Cahn	32	Global	0.0118	0.0123	0.0109
Allen–Cahn	32	POU-S (GN)	0.0130	0.0195	0.0209
Allen–Cahn	32	FT	0.0529	0.0483	0.0457
Burgers	0	Frozen	0.2750	0.2577	0.2454
Burgers	0	Global	0.0248	0.0277	0.0288
Burgers	0	POU-S (GN)	0.0289	0.0316	0.0336
Burgers	0	FT	0.0187	0.0215	0.0211
Burgers	32	Frozen	0.2750	0.2577	0.2454
Burgers	32	Global	0.0217	0.0240	0.0253
Burgers	32	POU-S (GN)	0.0245	0.0273	0.0298
Burgers	32	FT	0.0170	0.0178	0.0182
CDR	0	Frozen	0.1886	0.1937	0.2063
CDR	0	Global	0.1023	0.1065	0.0972
CDR	0	POU-S (GN)	0.1066	0.1198	0.1030
CDR	0	FT	0.0195	0.0198	0.0199
CDR	32	Frozen	0.1886	0.1937	0.2063
CDR	32	Global	0.0895	0.0927	0.0866
CDR	32	POU-S (GN)	0.0940	0.1040	0.0911
CDR	32	FT	0.0112	0.0113	0.0114

Table 13: **Separate first-use profiles for 350-step backbones on the RTX 2070 SUPER workstation.** One fixed validation instance for each of Allen–Cahn and Burgers, seed 11 and 32 sensors. Fresh-process first calls include backend initialization, weight transfer and compilation; warm values are medians of three subsequent calls. Disk reads and process imports are excluded. These six profiles are separate from the test-cohort averages.

PDE	Method	Steps	First call (s)	Warm (s)
Allen–Cahn	FT	5600	7.98	5.56
Allen–Cahn	Global	2	0.967	0.17
Allen–Cahn	POU-S (GN)	2	0.985	0.166
Burgers	FT	2800	5.23	2.77
Burgers	Global	8	1.12	0.327
Burgers	POU-S (GN)	8	1.14	0.321

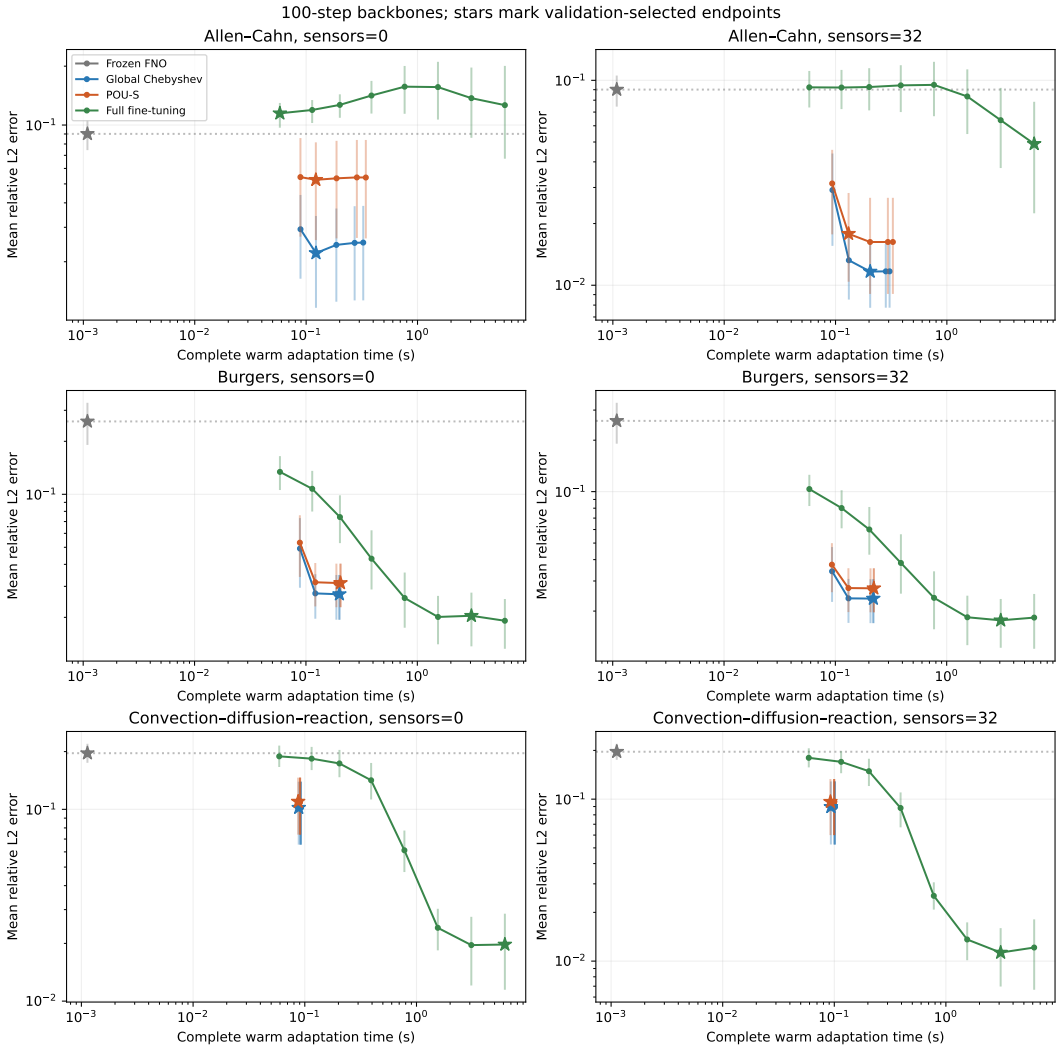


Figure 3: **Accuracy–time curves for the matched held-out FNO comparison.** The 100-step backbones, test instances and observations match Table 2. Vertical bars are 95% crossed-bootstrap intervals for the mean over instances and seeds; stars mark validation-selected budgets. Times include method setup and final prediction on the RTX 4090 host. The first-use profiles in Table 13 use the separate 350-step workstation setup.

Table 14: **Matched comparison with 350-step backbones.** Relative L_2 mean $\pm$ sample SD over 12 instances $\times$ three backbones per condition (crossed observations). Trainable coordinate counts are nominal. $m = 0$ uses PDE/IC/BC only; $m = 32$ adds sensors. Times are mean per-case median end-to-end warm seconds, including setup; parentheses give selected GN/Adam iteration budgets. The cohort and objective match Table 2; validation selects a separate recipe at each backbone strength. These times use the RTX 2070 SUPER workstation, not the RTX 4090 host used for Table 2.

PDE	m	Frozen	Global Cheb.	POU-S (GN)	FT
Trainable parameters		0	288	288	241K
Allen–Cahn	0	0.0135 ± 0.0037	0.0177 ± 0.0106	0.0150 ± 0.0088	0.0152 ± 0.0099
Time (s)		0.00122	0.163 (2)	0.160 (2)	5.53 (5600)
Allen–Cahn	32	0.0135 ± 0.0037	0.0116 ± 0.0067	0.0113 ± 0.0066	0.0116 ± 0.0058
Time (s)		0.00122	0.169 (2)	0.166 (2)	5.55 (5600)
Burgers	0	0.1058 ± 0.0644	0.0225 ± 0.0132	0.0250 ± 0.0141	0.0191 ± 0.0111
Time (s)		0.00122	0.259 (8)	0.264 (8)	5.57 (5600)
Burgers	32	0.1058 ± 0.0644	0.0199 ± 0.0117	0.0222 ± 0.0133	0.0181 ± 0.0108
Time (s)		0.00121	0.263 (8)	0.271 (8)	2.79 (2800)
CDR	0	0.0798 ± 0.0569	0.0631 ± 0.0411	0.0631 ± 0.0411	0.0202 ± 0.0158
Time (s)		0.00112	0.118 (1)	0.118 (1)	5.55 (5600)
CDR	32	0.0798 ± 0.0569	0.0494 ± 0.0373	0.0495 ± 0.0371	0.0126 ± 0.0096
Time (s)		0.00111	0.121 (1)	0.122 (1)	2.78 (2800)

Table 15: **Backbone-strength ablation on one Allen–Cahn instance.** Relative L_2 errors multiplied by 10^3, reported as mean $\pm$ sample SD over three paired backbone seeds. m is the sensor count; pretraining steps vary while adaptation budgets remain fixed.

Pretraining steps	m	Frozen	Global Cheb.	POU-S	FT
100	0	93.027 ± 20.072	1.556 ± 0.004	1.619 ± 0.133	4.282 ± 0.526
100	32	93.027 ± 20.072	1.083 ± 0.014	1.333 ± 0.200	3.583 ± 1.000
350	0	29.082 ± 1.956	1.552 ± 0.001	1.717 ± 0.303	2.174 ± 0.209
350	32	29.082 ± 1.956	1.155 ± 0.004	1.424 ± 0.381	1.615 ± 0.317

On this instance, undertraining increases mean frozen error by $3.2\times$, while final POU errors remain similar at both sensor counts (Table 15). Global Cheb. is similarly insensitive, while FT attains lower error from the stronger checkpoint under the fixed recipe. This observation does not extend uniformly to the full test cohort above. Pretraining changes error structure and feature conditioning; the single-instance comparison does not isolate an irreducible error floor.

C.2 Validation-selected FF-MLP refinement

This bounded fine-tuning sensitivity study follows the validation-only protocol in Appendix A.10; the selected recipes are listed in Table 10. The tables below report test outcomes and complete warm adaptation costs for the original test cohort.

Table 16: **FF-MLP test errors and complete warm adaptation costs.** Mean relative error $\pm$ sample SD over 12 instances $\times$ three seeds; other entries are arithmetic means. Time averages per-case medians of three synchronized repetitions. BC is the square root of the mean of the two side MSEs; the historical table instead averages side RMSEs.

PDE	Method	Rel-L_2	Time (s)	PDE	IC	BC
Allen–Cahn	Frozen	0.1862 ± 0.0257	–	4.09	0.0876	0.0692
Allen–Cahn	FT (validation-selected)	0.2065 ± 0.0234	0.209	4	0.0778	0.0634
Allen–Cahn	POU-D (sel. init.)	0.1083 ± 0.0551	0.458	0.152	0.0272	0.0215
Burgers	Frozen	0.4904 ± 0.2287	–	22.8	0.207	0.139
Burgers	FT (validation-selected)	0.4306 ± 0.2091	1.62	22.9	0.0669	0.0891
Burgers	POU-D (sel. init.)	0.1451 ± 0.1152	0.503	0.73	0.17	0.0798

Table 17: **Paired POU-D versus validation-selected FT.** Error ratios below one favor POU. Intervals use 5,000 crossed bootstrap resamples of the 12 instance IDs and three seeds, and remain exploratory with three seeds. The time ratio divides the FT mean by the POU mean; values above one indicate faster POU adaptation.

PDE	POU/FT error	95% interval	POU error wins	FT/POU time
Allen–Cahn	0.407	$[0.234, 0.635]$	36/36	0.456
Burgers	0.294	$[0.241, 0.358]$	36/36	3.22

The validation-selected recipe changes mean FT test error from the original fixed 400-step result as follows: Allen–Cahn 0.209 to 0.207; Burgers 0.44 to 0.431. The POU solution errors reproduce the archived compact-window results. The comparison tests FT recipe sensitivity while retaining the different fitting grids and regularizers; the shared-objective FNO study remains the principal controlled comparison.

C.3 Solution and constraint diagnostics

The following tables separate solution error from each fitted constraint. Residuals use the same discrete operator, points and masks for all methods within a comparison. FNO and FF-MLP POU residuals are fitting-grid diagnostics; they do not certify a continuous-domain PDE error bound. A decrease of the combined objective need not decrease every component, and lower residual need not imply lower solution error. Dashes denote inapplicable constraints.

The Burgers PDE-RMSE change from 0.61 to 0.73 compares the archived Gaussian-window adapter with the compact-window rerun. It is not the change from the frozen field to the adapted field; Table 19 gives that paired comparison. The support radius was fixed before the compact rerun.

Table 18: **FNO solution and constraint errors for Table 2.** Arithmetic means over 36 case/seed pairs. "Unobs." excludes initial/boundary traces and sensors. Boundary RMSE is the square root of the mean of the two side MSEs. No interface constraint applies.

PDE	m	Method	Rel-L_2	Unobs.	PDE	IC	BC	Sensors
Allen–Cahn	0	Frozen	0.0901	0.0835	0.572	0.0805	0.089	–
Allen–Cahn	0	Global	0.0222	0.0223	0.0378	0.000693	$9.03e-06$	–
Allen–Cahn	0	POU-S (GN)	0.0525	0.0523	0.0484	0.00212	0.00539	–
Allen–Cahn	0	FT	0.115	0.114	0.433	0.0322	0.0606	–
Allen–Cahn	32	Frozen	0.0901	0.0833	0.572	0.0805	0.089	0.0716
Allen–Cahn	32	Global	0.0117	0.0119	0.0248	0.000554	$9.81e-06$	0.00116
Allen–Cahn	32	POU-S (GN)	0.0178	0.0176	0.0567	0.00173	0.00664	0.00199
Allen–Cahn	32	FT	0.049	0.0501	0.0138	0.000479	0.00134	0.000853
Burgers	0	Frozen	0.259	0.285	1.68	0.119	0.0886	–
Burgers	0	Global	0.0271	0.0291	0.0767	0.00444	0.012	–
Burgers	0	POU-S (GN)	0.0313	0.0327	0.0938	0.00586	0.0176	–
Burgers	0	FT	0.0204	0.0232	0.00713	0.00201	0.00161	–
Burgers	32	Frozen	0.259	0.284	1.68	0.119	0.0886	0.118
Burgers	32	Global	0.0237	0.0254	0.0815	0.00498	0.0115	0.00509
Burgers	32	POU-S (GN)	0.0272	0.0283	0.1	0.00677	0.0159	0.00614
Burgers	32	FT	0.0177	0.0202	0.0174	0.00242	0.00366	0.00252
CDR	0	Frozen	0.196	0.2	0.923	0.125	0.153	–
CDR	0	Global	0.102	0.106	0.247	0.0138	0.0587	–
CDR	0	POU-S (GN)	0.11	0.113	0.269	0.0194	0.0755	–
CDR	0	FT	0.0197	0.0209	0.00882	0.000416	0.00322	–
CDR	32	Frozen	0.196	0.2	0.923	0.125	0.153	0.154
CDR	32	Global	0.0896	0.0941	0.296	0.0153	0.0566	0.0429
CDR	32	POU-S (GN)	0.0964	0.0998	0.322	0.021	0.0736	0.0486
CDR	32	FT	0.0113	0.0119	0.0355	0.0014	0.00695	0.00301

Table 19: **Archived FF-MLP same-instance refinement diagnostics.** Means over 12 instances and three seeds. BC is the mean of the two side RMSEs, as in the archive. The global row matches coefficient count only; its basis and one-shot fitting differ from the conditioned, polished POU fit. FT uses the original 400-step recipe.

PDE	Method	Rel-L_2	PDE	IC	BC
Allen–Cahn	Frozen	0.186	4.09	0.0876	0.0691
Allen–Cahn	Global (one-shot)	0.652	13	0.313	0.247
Allen–Cahn	POU-D (sel. init.)	0.108	0.152	0.0272	0.0208
Allen–Cahn	FT (fixed budget)	0.209	3.91	0.0717	0.0401
Burgers	Frozen	0.49	22.8	0.207	0.138
Burgers	Global (one-shot)	9.04	247	2.3	0.437
Burgers	POU-D (sel. init.)	0.145	0.73	0.17	0.0797
Burgers	FT (fixed budget)	0.44	22.9	0.0692	0.0996

Table 20: **CDR geometry comparison.** Twelve held-out manufactured instances and three backbone seeds (11, 17, 23); physics/IC/BC fitting without sensors. BC uses the archived mean of side RMSEs. FT retains its fixed 180-step budget; it is not the validation-selected FNO comparison.

Method	Rel-L_2 mean $\pm$ SD	PDE	IC	BC
Frozen	0.0983 ± 0.0508	0.804	0.0382	0.076
Spatial POU-S	0.1368 ± 0.0909	0.138	0.0885	0.0984
Characteristic POU-S	0.0464 ± 0.0586	0.00842	0.00385	0.0314
FT (fixed budget)	0.0421 ± 0.0371	0.0731	0.005	0.0288

Table 21: **Material and interface constraint diagnostics for Table 4.** Deterministic cases, not seed averages. The κ rows omit interface constraints from fitting; value and flux errors are evaluation diagnostics. Circle/sunflower rows include them. Jump errors subtract the prescribed jumps using the archived offset traces and derivative convention; finite trace offsets contribute discretization error. PDE magnitudes are not comparable across conductivity scales.

Case	Method	Rel-L_2	Bulk PDE	Boundary	Value jump	Flux jump
$k_R = 2$	Frozen	0.107	1.86	2.30×10^{-18}	0.000204	0.2
$k_R = 2$	Global	0.101	1.52	2.50×10^{-6}	0.00026	0.256
$k_R = 2$	POU-S	0.09	1.41	2.30×10^{-18}	0.000338	0.333
$k_R = 2$	POU-C	0.0565	1.23	4.67×10^{-5}	0.000385	0.381
$k_R = 2$	POU-D	0.0624	0.99	2.47×10^{-5}	0.000384	0.379
$k_R = 5$	Frozen	0.118	4.03	2.45×10^{-18}	0.000204	0.808
$k_R = 5$	Global	0.115	2.68	6.67×10^{-6}	0.00022	0.871
$k_R = 5$	POU-S	0.113	2.41	3.46×10^{-5}	0.000204	0.808
$k_R = 5$	POU-C	0.102	1.67	5.55×10^{-5}	0.000177	0.698
$k_R = 5$	POU-D	0.106	1.67	2.08×10^{-5}	0.000182	0.716
$k_R = 10$	Frozen	0.128	9.53	2.62×10^{-18}	0.000204	1.82
$k_R = 10$	Global	0.124	6.44	8.70×10^{-6}	0.000203	1.81
$k_R = 10$	POU-S	0.119	5.78	4.95×10^{-5}	0.000204	1.82
$k_R = 10$	POU-C	0.126	2.26	7.25×10^{-5}	0.000121	1.07
$k_R = 10$	POU-D	0.118	2.38	2.46×10^{-5}	9.57×10^{-5}	0.842
Circle	Frozen	0.587	$3.86e+03$	2.65×10^{-16}	0.0292	$2.33e+03$
Circle	Global	0.56	179	4.27×10^{-6}	0.00273	157
Circle	POU-C	0.433	$1.03e+03$	0.00064	0.00305	405
Circle	Subdomain POU-C	0.479	325	3.29×10^{-6}	0.00159	52.9
Sunflower	Frozen	0.878	53.1	3.44×10^{-16}	0.31	23.2
Sunflower	Global	1.68	96.3	0.00042	0.27	49.1
Sunflower	POU-C	3.58	114	0.00209	0.254	54.9
Sunflower	Subdomain POU-C	0.694	4.63	0.000124	0.103	1.54

Table 22: **Singular-repair constraint diagnostics.** Means over four coarse resolutions and three sensor-placement seeds, with $m = 32$. Unchanged fits were replayed to recover unsaved boundary/sensor errors and reproduce archived solution/PDE errors. PDE RMSE here is the raw grid residual, including the discrete source for point-source cases; it differs from the excess/no-source diagnostics used elsewhere. The source flux column reports absolute annular flux error relative to unit source strength.

Case	Method	Global rel-L_2	PDE	Boundary	Sensors	Flux
Corner	Frozen	0.0318	25.7	0	0.0414	–
Corner	Global	0.0563	25.7	0.013	0.0108	–
Corner	POU-S	0.0199	24.8	0.00756	0.00258	–
Corner	Enriched POU	0.0186	27.9	0.00353	0.00221	–
Crack	Frozen	0.0346	5.73	0	0.0579	–
Crack	Global	0.0361	7.91	0	0.029	–
Crack	POU-S	0.0168	5.72	0	0.00558	–
Crack	Enriched POU	0.016	5.23	0	0.0033	–
Source	Frozen	0.0637	154	0	0.061	0.0594
Source	Global	0.0687	154	0	0.0598	0.0663
Source	POU-S	0.056	153	0	0.0202	0.134
Source	Source-aware POU	0.039	166	0	0.00619	0.00653

C.4 HISTORICAL FIXED-BUDGET COMPARISON

Table 23 preserves the earlier compact-window comparison. Its FNO Allen–Cahn/Burgers rows repeat one instance across three seeds; FF-MLP and CDR use 12 instances and three seeds. The FNO global basis lacks frozen-field conditioning, and summed least-squares blocks differ from FT block means. These rows establish the original fixed-budget behavior, while Table 2 supplies the matched multi-instance comparison.

The archived 576-coefficient FF-MLP Burgers global control uses unconditioned spatial Fourier columns through frequency index 144, with constant/linear time factors and a one-shot ridge fit (10^{-7}). On the 64-point spatial grid, high-frequency values alias while their continuous derivatives remain frequency-dependent. POU instead uses $[1, g, g^2]$ conditioning and nonlinear coefficient polish. These representation and solver differences prevent attributing the large global-control error to locality alone.

On the unobserved interior points of the historical single-instance FNO $m = 32$ rows, rel-L2 is 0.0284 (frozen), 0.0168 (global), 0.0020 (POU-S), and 0.0095 (fine-tuning) for Allen–Cahn, and 0.285, 0.375, 0.108, and 0.069 for Burgers, so the ranking matches the historical Table 23.

Table 23: **Historical fixed-budget compact-window results.** FF-MLP refines a model pretrained on the same instance; its POU-D row uses selection-dependent initialization. Control matches coefficient count only, except CDR where it denotes standard spatial POU. Archived costs use pipeline-specific timers, exclude common frozen-field evaluation, and are not end-to-end speed comparisons.

Backbone	PDE	Signal	Mode	Frozen	Control	POU	FT	POU cost	FT cost
FF-MLP	Burgers	physics	POU-D	0.490	9.040	**0.145**	0.440	576/1.96s	4.8K/3.62s
FF-MLP	Allen–Cahn	physics	POU-D	0.186	0.652	**0.108**	0.209	576/1.78s	4.8K/3.54s
FNO	Allen–Cahn	$m = 32$ sensors	POU-S	0.0296	0.0169	**0.00201**	0.00939	288/21ms	241K/1.75s
FNO	Burgers	$m = 32$ sensors	POU-S	0.273	0.369	0.104	**0.0658**	288/25ms	241K/1.74s
FNO	CDR	physics	POU-S	0.0983	0.1368	0.0464	**0.0421**	288/61ms	241K/1.51s

C.5 HISTORICAL FNO SELECTION AND CAPACITY ABLATIONS

These historical diagnostics use one held-out instance per PDE and average over three backbone seeds, matching the single-instance FNO rows of Table 23. Coordinate fitting uses a one-shot linearized ridge solve. They describe selection and capacity effects in that setting; they are not ablations of the multi-instance, validation-selected Gauss–Newton comparison in Table 2.

Table 24: **Historical FNO selection and capacity diagnostics.** Means over three backbone seeds on one held-out instance per PDE, using one-shot linearized ridge fitting. Dense target values are excluded from coefficient fitting; the oracle variant uses them for packet selection only. The reference setting is residual-selected POU-S with $K = 6$ and degree 3.

Ablation	Equation / regime	Compared variants	Params	rel-L2 ↓
Packet selection	Allen–Cahn FNO, $m = 32$	residual / random / sensor-aware / oracle	288 (all variants)	0.0020 / 0.0065 / 0.0062 / 0.0075
Packet selection	Burgers FNO, $m = 32$	residual / random / sensor-aware / oracle	288 (all variants)	0.1037 / 0.1095 / 0.1045 / 0.1064
Active-packet budget	Allen–Cahn FNO, $m = 32$, degree 3	$K = 2$ / 4 / 6 / 8	96 / 192 / 288 / 384	0.0179 / 0.0136 / 0.0020 / 0.0015
Local basis degree	Allen–Cahn FNO, $m = 32$, $K = 6$	degree 1 / 2 / 3 / 4	72 / 162 / 288 / 450	0.0290 / 0.0047 / 0.0020 / 0.0016

Residual selection has the lowest mean error in these single-instance diagnostics, with sparse sensors included in fitting. Oracle selection uses dense target error with the same fitting objective, so it is not an error floor; sensor-aware reselection can move packets away from useful residual supports. The Burgers selection rules differ by at most 6%, consistent with a broad transport error. More packets and higher degree improve this Allen–Cahn instance. The primary multi-instance comparison fixes

$K = 6$ and degree 3 and selects fitting budgets on separate validation instances; these historical sweeps do not establish selection or capacity trends across that cohort.

C.6 Support radius and locality in the learned-solver rows

In the historical frozen-solver rows of Table 23, the six selected supports of radius $r_p = 0.72$ (12 centers spaced $h = 1/6$ on $[-1, 1]$) jointly cover the whole spatial domain, so the correction is not local there. To measure the effect of locality, we repeat the POU rows with radii $r_p = h$ and $r_p = 2h$, fixed before running, keeping every other setting unchanged. Table 25 reports rel-L2 and the fraction of the domain covered by the selected supports; outside this union the frozen field is unchanged. For CDR the coverage is measured in the characteristic coordinate $\zeta = x - vt$. The FF-MLP POU-D rows refit all packets, with stronger ridge on unselected coefficients during initialization, so their correction covers the domain at every radius.

Table 25: **Support radius in the frozen-solver rows.** Entries give rel-L2 and, for POU-S, the fraction of the domain covered by the selected supports. Frozen and fine-tuning values do not depend on the radius.

Row	Mode	$r_p = h$	$r_p = 2h$	$r_p = 0.72$	Frozen	FT
FNO Allen–Cahn, $m = 32$	POU-S	0.0150 (78%)	0.0114 (93%)	0.0020 (100%)	0.0296	0.0094
FNO Burgers, $m = 32$	POU-S	0.129 (67%)	0.114 (96%)	0.104 (100%)	0.273	0.066
FNO CDR, characteristic	POU-S	0.078 (73%)	0.060 (93%)	0.046 (100%)	0.098	0.042
FF-MLP Burgers	POU-D	0.125	0.129	0.145	0.490	0.440
FF-MLP Allen–Cahn	POU-D	0.070	0.078	0.108	0.186	0.209

With local supports ($r_p = h$), POU-S leaves 22–33% of the domain unchanged and still improves on the frozen model in every FNO Allen–Cahn and Burgers run and in 33 of 36 CDR cases; it also improves on the historical global controls of Table 23 and, for CDR, on standard spatial packets (0.102). It no longer matches fine-tuning, which is better in every run. Wider supports improve accuracy for these broad learned-solver errors, whereas the localized errors of Sections 4.2 and 4.3 are repaired with compact supports that leave a large part of the domain unchanged. The FF-MLP POU-D rows improve as the radius decreases, which reflects the conditioning and capacity of the full-scaffold fit rather than locality.

C.7 Damaged-repair budget robustness

The local-error table fixes the smooth adapter budget at 40 coefficients. To check that the comparison is not tied to this single budget, Table 26 repeats the crack-tip and point-source repair settings with $m = 32$ sensors and smooth correction budgets of 20, 40, and 80 coefficients. The global Chebyshev control uses the same number of coefficients as smooth POU-S; the problem-aware POU variant adds one singular or Green enrichment coefficient per active packet.

Table 26: **Damaged-repair budget robustness.** Entries report local rel-L2, with global rel-L2 in parentheses. Dense target fields are used only for evaluation.

Error structure	Coeffs.	Global	POU-S	Geometry/source-aware POU
Crack tip	20	0.0888 (0.0476)	0.0526 (0.0203)	**0.0502** (0.0201)
Crack tip	40	0.0782 (0.0361)	0.0278 (0.0168)	**0.0167** (0.0160)
Crack tip	80	0.0830 (0.0331)	0.0346 (0.0141)	**0.0122** (0.0123)
Point source	20	0.0810 (0.0680)	0.1088 (0.0958)	**0.0378** (0.0449)
Point source	40	0.0811 (0.0687)	0.0564 (0.0560)	**0.0298** (0.0390)
Point source	80	0.0816 (0.0745)	0.0528 (0.0572)	**0.0252** (0.0480)

The crack-tip result is robust to the coefficient budget: smooth POU-S improves over the matched global correction at all three budgets, and the singular-enriched packet improves further. The point-source result shows the scope boundary more sharply. For the point-source problem, smooth POU-S increases local error relative to the frozen field in $11/12$, $2/12$, and $2/12$ cases at the 20-, 40-, and

80-coefficient budgets, respectively. The source-aware variant reduces local error in all 12 cases at each budget, although global and outside-region errors can still increase.

C.8 POROUS TOPOLOGY-UPDATE RESULTS

This topology-update study supplements Section 4.3; its equations and paired-field setup are given in Appendix B.3.

Porous void-pair local topology update. The supervised porous heat diagnostic fits the dense residual $u_B - u_A$ and is therefore retained only as a locality/capacity check. A deployment version of the same paired topology-change setting uses configuration B only through the new geometry, coefficient map, source map, boundary traces, and optional sparse measurements, while the dense B solution is withheld until evaluation. The result is positive but narrower than the supervised diagnostic. With $m = 128$ sparse measurements, residual-selected POU-S improves on the frozen anchor and the matched global Chebyshev control at the same 160-coefficient budget; selecting packets by overlap with the known changed-topology band instead gives a small further gain. At $m = 32$, both POU-S variants worsen the physical-domain and void-inclusive errors. At $m = 64$, physical-domain error improves modestly, while the void-inclusive error increases; at $m = 128$, both metrics improve. We therefore use this benchmark as appendix evidence for sparse-anchor topology repair and as a scope warning for physics-only material/topology claims.

Table 27: **Sparse porous void-pair deployment.** Values are means over two topology pairs and three sensor-placement seeds. The table reports outer-domain rel-L2, with domain-B rel-L2 in parentheses. Global Chebyshev, residual-selected POU-S, and topology-selected POU-S each use 160 trainable coefficients; both dense POU-D variants use 576 coefficients and are discussed in the text. Dense target fields are not used for fitting. The domain-B error (in parentheses) is the physical-domain metric; the outer-domain error also counts void cells.

m	Frozen	Global Cheb.	Residual POU-S	Topology POU-S
32	0.2117 (0.1482)	0.2192 (0.1481)	0.2384 (0.1575)	0.2299 (0.1537)
64	0.2117 (0.1482)	0.2128 (0.1467)	0.2230 (0.1443)	0.2163 (0.1423)
128	0.2117 (0.1482)	0.2101 (0.1367)	0.1956 (0.1180)	**0.1911** (**0.1164**)

At $m = 128$, two POU-D variants fit all 36 packets (576 coefficients). Topology-initialized POU-D uses the selected POU-S solution as a coefficient prior with different selected/unselected ridge weights; it gives domain-B rel-L2 0.1152 and boundary RMSE 0.216. Uniformly regularized POU-D gives 0.1159 and 0.189, respectively. Both improve domain-B error over topology-selected POU-S, but neither improves its void-inclusive error. They are dense-capacity checks; fitting all packets distinguishes them from the limited coordination channel of POU-C. Figure 4 visualizes the resulting adapted solution and active packet centers across the modified void geometry.

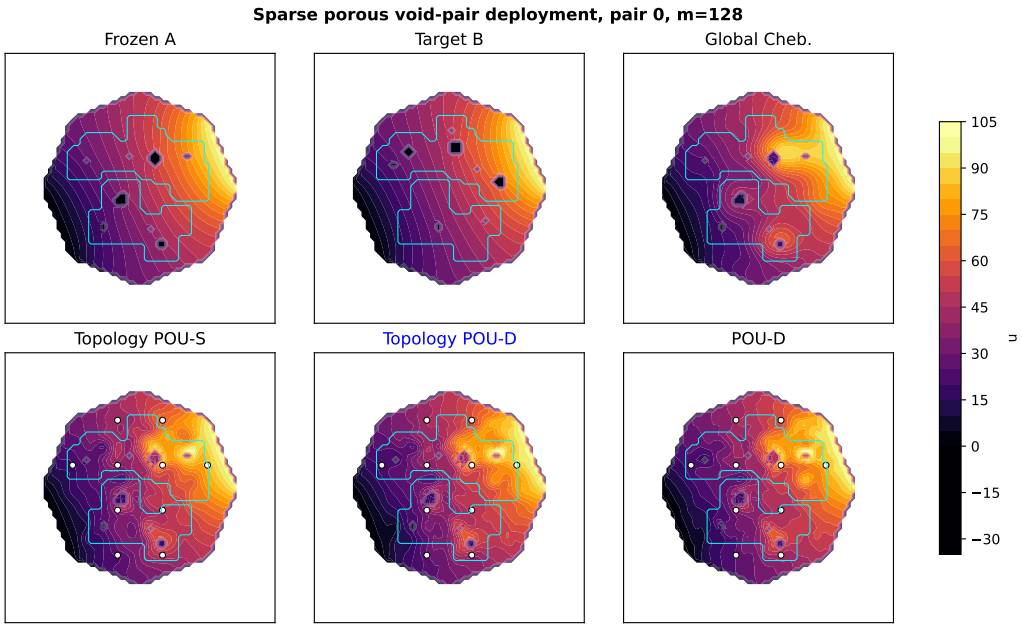


Figure 4: **Representative porous heat void-pair deployment update.** The target configuration changes the internal void geometry while preserving the outer body. The sparse-topology deployment setting fits from the new operator, boundary traces, sparse measurements, and a weak preservation prior on unchanged solid cells rather than the dense target update; the figure overlays the topology-change band and selected POU centers. "Topology POU-D denotes the topology-initialized, trust-regularized dense fit; "POU-D uses uniform regularization. White markers show the topology-selection set; both dense panels fit all 36 packets.

D SUPPLEMENTARY LOCALITY AND CAPACITY DIAGNOSTICS

These controlled studies examine locality, correction capacity, and backbone quality. They complement the deployment experiments, with supervised fitting signals and diagnostic assumptions stated for each study.

D.1 CONTROLLED LOCALITY DIAGNOSTICS

Fourier-feature MLP setup. This is a synthetic one-dimensional diagnostic on $x \in [0, 1]$, not a PDE solve. A Fourier-feature MLP is first trained on a smooth generic base field and frozen. The target then adds a compact oscillatory local defect. This diagnostic tests whether an external coordinate-space POU adapter can correct a black-box implicit model more locally than LoRA, full fine-tuning, or a budget-matched global coordinate correction.

LoRA versus coordinate-space POU adapters. The scalar-output FF-MLP has a 64×1 final weight matrix. A LoRA update AB, with $A \in \mathbb{R}^{64\times r}$ and $B \in \mathbb{R}^{r\times 1}$, has rank at most one for every nominal r. Thus the archived final-layer settings $r = 2, 4, 8$ change factorization, initialization and the $1/r$ scaling, not representational rank; their parameter counts are $65r$. The meaningful final-layer comparison uses $r = 1$ (65 parameters); all-layer ranks 1 and 2 also adapt hidden matrices and are distinct capacity settings. This supervised diagnostic examines local approximation and spatial addressability, not broad superiority over optimized weight-space adaptation. The all-layer rank-2 result is shown in Figure 5; higher nominal final-layer ranks are retained only as redundant-factorization diagnostics in the experiment archive.

Support-informed global control. Table 29 gives a masked global correction the same selected-window envelope $m = \sum_{p\in S} \bar{w}_p$, features $[1, g, g^2]$, and 72 coefficients as selected POU. This

Table 28: **Nonredundant LoRA settings in the archived supervised local-defect diagnostic.** One manufactured instance, seed 7. Final-layer rank 1 spans every possible scalar-output weight update; all-layer ranks also affect hidden matrices. Errors are diagnostic relative L_2 values, not PDE residuals or multi-instance estimates.

Method	Parameters	Global error	Defect error	Outside error
Final-layer LoRA, rank 1	65	0.03747	0.06200	0.03198
All-layer LoRA, rank 1	294	0.00335	0.00525	0.00294
All-layer LoRA, rank 2	588	0.00249	0.00434	0.00205
Selected conditioned POU	72	0.00614	0.01588	0.00205
Dense conditioned POU	432	0.00128	0.00324	0.00050
Full fine-tuning	6657	0.00139	0.00189	0.00130

separate supervised run uses one manufactured instance (seed 7), with 301 fitting points drawn at stride four from a 1,200-point grid, including the final endpoint. The two packets are selected by windowed target-error energy on the fitting points; both corrections use ridge 10^{-6}. The global space has 24 shared coordinate columns (constant, linear, and Fourier terms), whereas POU has 12 degree-11 Chebyshev terms in each of two local packets. Thus the comparison tests the specified representations with equal location information, rather than isolating locality from every basis choice. POU has lower error in this diagnostic; the supplied mask alone does not reproduce its accuracy.

Table 29: **Same-mask comparison in the supervised local-defect diagnostic.** Both methods use $[1, g, g^2]$ and 72 coefficients. Relative L_2 errors use the full 1,200-point grid, the true defect region $|x - 0.68| \leq 0.12$, or its complement. The selected support extends beyond the defect; outside-defect error therefore differs from error outside the correction support. Both corrections vanish outside their common support.

Method	Coefficients	Global error	Defect error	Outside error
Masked global	72	0.04290	0.11163	0.01346
Selected POU	72	0.00614	0.01588	0.00205

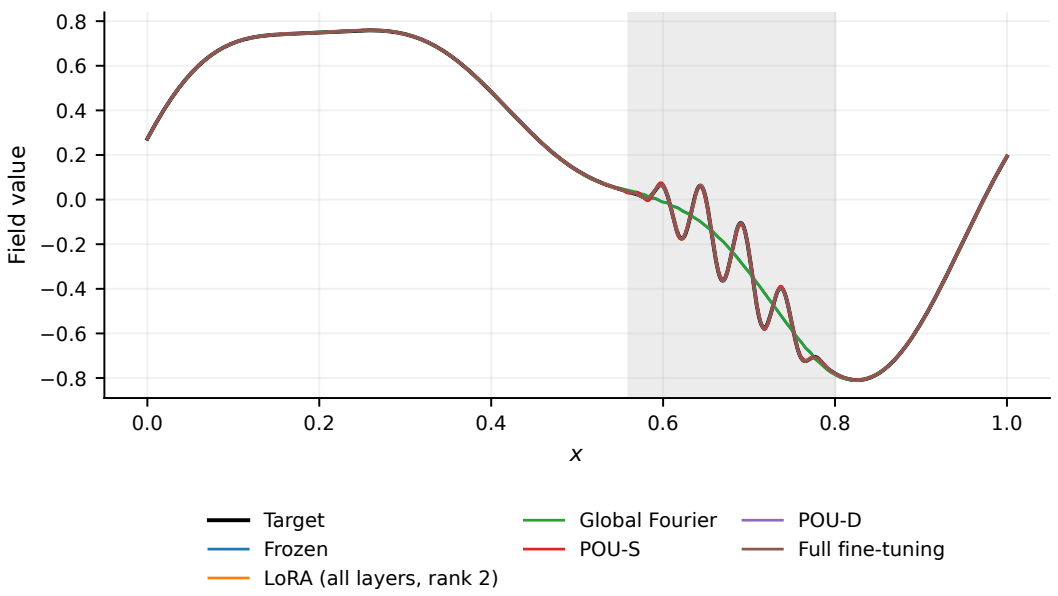


Figure 5: **Controlled local-defect diagnostic with a Fourier-feature MLP backbone.** This supervised non-PDE diagnostic uses a one-dimensional base field plus a compact oscillatory defect (shaded region). POU-S and POU-D denote selected and all-packet corrections conditioned on the frozen field. The adapter comparison tests local approximation capacity.

Synthetic local-defect diagnostics. Earlier manufactured Poisson bump, coarse-patch, and ringing corruptions remain useful as guardrail diagnostics: they test whether preservation-aware objectives can accept, damp, or reject a proposed local correction. They are not used as the main localized/singular-error evidence because the classical singular benchmarks in Appendix B.2 provide a more principled numerical-repair setting.

D.2 Manufactured RBF local-damage diagnostics

Setup and fitting signals. An anisotropic RBF model with learned centers, widths, weights, and rotations is trained on a smooth two-dimensional manufactured target, frozen, and given a compact output-level defect. The supervised variant fits $u_{\text{target}} - u_{\text{damaged}}$ directly. The calibrated Poisson variant uses exact-autodiff residuals $r_{\text{pre}} = -\Delta g_{\text{pre}} - f$ and $r_{\text{damaged}} = -\Delta g_{\text{damaged}} - f$: selection and fitting use $r_{\text{base}} = r_{\text{damaged}} - r_{\text{pre}}$, with Dirichlet data $b_{\text{base}} = g_{\text{damaged}} - u_{\text{target}}$ on the boundary. All adapters in Figure 6 fit this calibrated objective, which targets restoration of the pre-damage residual and requires access to it. Interior target values serve backbone pretraining and diagnostic error reporting, but do not enter the Poisson adapter fit.

Coordinated correction in the Poisson diagnostic. The correction family and least-squares objective are defined in Appendix A.8. A separate coordinated Poisson diagnostic uses a degree-3 tensor-product global Chebyshev basis B_R with $(3+1)^2 = 16$ coefficients. Four selected degree-5 POU packets use $4(5+1)^2 = 144$ coefficients, so this relaxed variant uses $144 + 16 = 160$. It uses the same calibrated residual and a field penalty limiting correction magnitude. Exploratory fits with a frozen RBF relaxation basis yielded higher solution errors than the Chebyshev variant for the tested trained backbones. Figure 6 shows global, selected, and all-packet corrections without the coordinated variant.

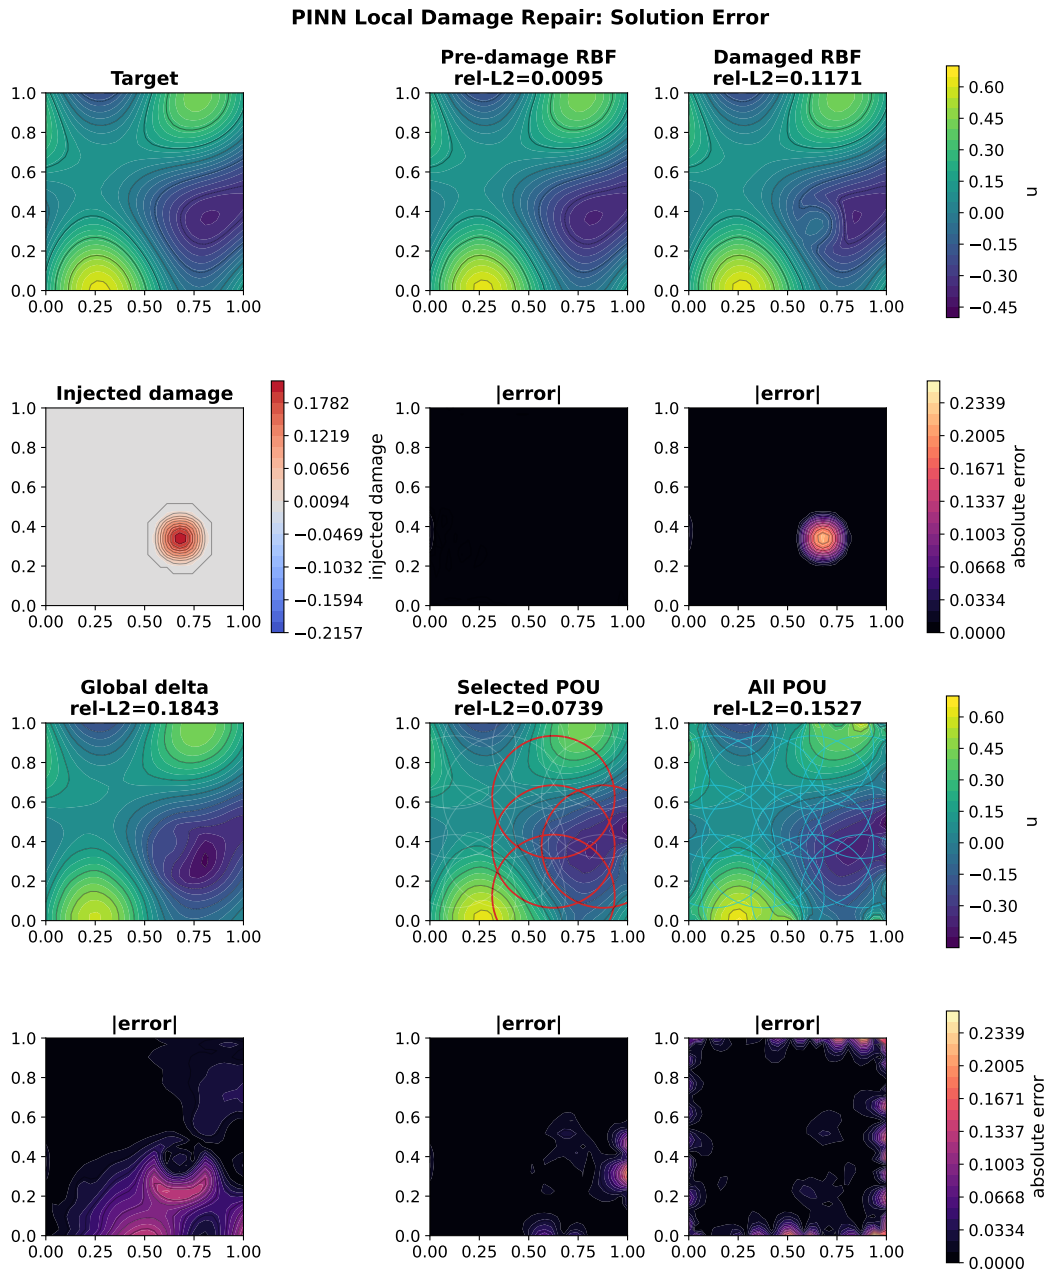


Figure 6: **Calibrated Poisson defect repair with an RBF backbone.** After freezing the backbone and injecting a compact defect, all displayed adapters fit the residual increase relative to the saved pre-damage field, with Dirichlet boundary data. This diagnostic assumes access to the pre-damage residual; interior target values are not used for adapter fitting.

D.3 Backbone quality and relaxation guidance

This supervised diagnostic varies an anisotropic RBF backbone's training budget, freezes it, and adds a local defect, a local defect with smooth global drift, or two separated defects. Selection and ridge fitting use the dense correction target $u_{\text{target}} - g_{\text{damaged}}$. Figure 7 and Table 30 compare selected POU with and without global RBF relaxation: local-plus-global repair error falls by 55.0%–79.6% across budgets, with smaller gains for strong backbones with mostly local defects.

Table 30: **Supervised backbone-quality diagnostic, local-plus-global regime.** Repair error is rel-L2 divided by that of the damaged frozen field; relaxation gain is its relative reduction. Selection and fitting use dense target values. Selected POU has 144 coefficients; adding 36 frozen RBF coefficients gives 180.

Backbone steps	Selected POU only	+ Global relaxation	Relaxation gain
20	0.644	0.290	55.0%
80	0.374	0.112	70.1%
250	0.372	0.088	76.4%
1200	0.369	0.075	79.6%

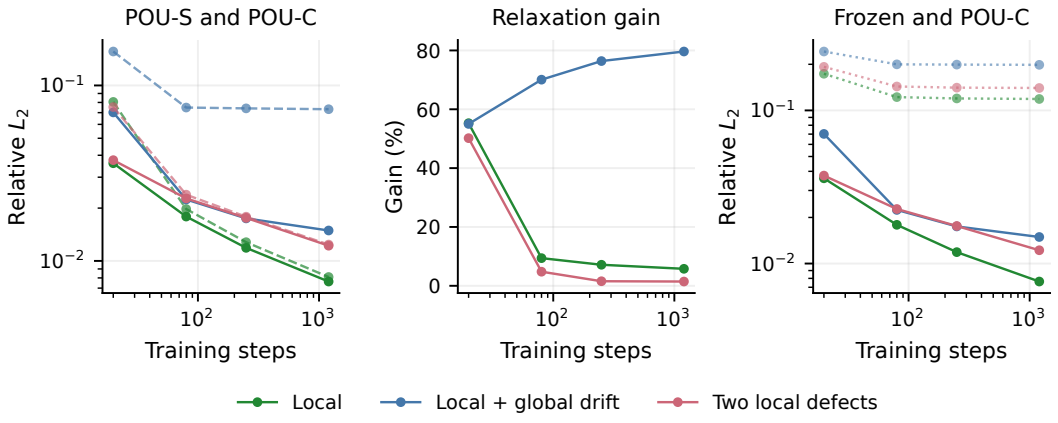


Figure 7: **Supervised backbone-quality and relaxation diagnostic.** Dashed curves are POU-S; solid curves are POU-C with global RBF relaxation; dotted curves show the damaged frozen field. Selection and fitting use dense target values. Gains are modest for a strong backbone with local errors and larger with smooth global drift or a weak backbone.

D.4 Multiscale basis and conditioning diagnostic

Supervised 1D setup. The target solves $-(au')' = 1$ on $[0,1]$ with $u(0) = u(1) = 0$ and $a(x) = 2 + 0.9\sin(2\pi x/0.03)$, by finite differences on 2049 points. The frozen coarse field is the homogenized solution $x(1-x)/(2\bar{a})$, with $\bar{a}$ the harmonic mean of a. The adapter is fit by ridge regression to the dense correction target $u_{\mathrm{ref}} - g$ at every second grid point.

Basis and conditioning comparison. On the fixed two-scale scaffold and frozen coarse field, all variants reduce error by 38–39% (Table 31). Unconditioned bases are nearly tied at 32 coefficients. Conditioning improves both families while adding coefficients; Fourier is slightly better at 64, and Chebyshev $[1, g, g^2]$ gives the lowest reported error.

Table 31: **Basis and conditioning comparison on a 1D multiscale elliptic target.** The scaffold and frozen field are fixed; the basis, features, and coefficient count vary.

Local basis	Params	Rel-L2	Improvement
Frozen homogenized coarse	0	0.00817639	0.00%
Chebyshev POU $[1]$	32	0.0050682	38.01%
Chebyshev POU $[1, g]$	64	0.0050322	38.45%
Chebyshev POU $[1, g, g^2]$	96	0.0049972	38.88%
Local Fourier POU	32	0.0050702	37.99%
Coarse-conditioned Fourier POU	64	0.0050011	38.84%

E RBF ADAPTER CONSTRUCTION AND SOLVE SCALING

This supervised scaling diagnostic uses the anisotropic-RBF local-damage setup (Appendix D.2) with a 6×6 candidate scaffold, separating selected from all-packet corrections. The backbone is trained once and frozen; selection and fitting use the dense correction target. On the Threadripper 2950X workstation, host-side NumPy builds the basis and solves the ridge system in double precision on a 48×48 grid. Reported times include only basis construction and solving, excluding selection, backbone training, and prediction evaluation. End-to-end FNO and FF-MLP timing protocols are in Appendices A.9 and A.10, with results in Appendices C.1 and C.2.

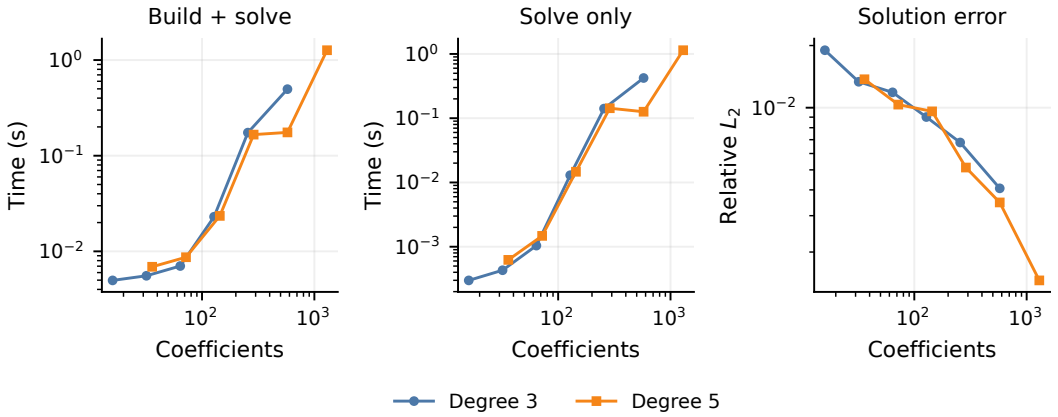


Figure 8: **Adapter timing scalability on the RBF local-damage diagnostic.** Timing is reported for selected POU packets with $K = 1, 2, 4, 8, 16$ active packets and the dense POU-D case (all 36 packets) on a 6×6 scaffold. The curves compare build-plus-solve time, solve time, and solution error as the active coefficient count increases.

Table 32: **Median adapter solve timings over three repeats.** The frozen backbone training time is excluded. Build time is POU basis construction; solve time is the ridge linear solve. Each timing column is aggregated separately: total is the median of the per-repeat build-plus-solve sums, which need not equal the sum of the build and solve medians.

degree	active patches	coeffs	rel-L2	build (s)	solve (s)	total (s)
3	1	16	0.01898	0.004671	0.0002991	0.00497
3	4	64	0.01187	0.005989	0.001041	0.007027
3	16	256	0.006785	0.03323	0.1411	0.1744
3	all	576	0.004069	0.07395	0.4224	0.4964
5	1	36	0.01373	0.006274	0.0006241	0.006899
5	4	144	0.009598	0.01026	0.01468	0.02349
5	16	576	0.003477	0.04916	0.1261	0.1753
5	all	1296	0.001458	0.1014	1.143	1.267